Palo Alto Networks®

PAN-OS® Getting Started Guide
PAN-OS 6.0

Contact Information

Corporate Headquarters:
Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054-1211

http://www.paloaltonetworks.com/contact/contact/

About this Guide

This Getting Started Guide takes you through the initial configuration and basic setup on your Palo Alto Networks firewalls. This guide takes over after you have completed rack mounting your hardware-based firewall or have created your virtual firewall; it is intended for administrators who want the basic framework to quickly set up the firewall as a security gateway.

For additional information, refer to the following sources:

- For information on the additional capabilities and for instructions on configuring the features on the firewall, go to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base, complete documentation set, discussion forums, and videos, go to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, or to manage your account or devices, go to https://support.paloaltonetworks.com.
- For the latest release notes, go to the Software Updates page at https://support.paloaltonetworks.com/Updates/SoftwareUpdates.

This guide provides procedures for configuring the firewall using the web interface on the device. It does not provide procedures for deploying firewalls using Panorama. For more information on using Panorama, refer to the Panorama Administrator's Guide.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2014 Palo Alto Networks. All rights reserved.

Palo Alto Networks, PAN-OS, and Panorama are trademarks of Palo Alto Networks, Inc. All other trademarks are the property of their respective owners.

P/N 810-000137-00B

Revision Date: October 15, 2014

Table of Contents

Integrate the Firewall into Your Management Network 1

Set Up Management Access to the Firewall 2
- Determine Your Management Strategy 2
- Perform Initial Configuration 2
- Set Up Network Access for External Services 5

Activate Firewall Services 9
- Register With Palo Alto Networks 9
- Activate Licenses 9
- Manage Content Updates 11
- Install Software Updates 13

Add Firewall Administrators 15
- Administrative Roles 15
- Administrative Authentication 15
- Create an Administrative Account 16

Monitor the Firewall 19
- Monitor Applications and Threats 19
- View Local Log Data 20
- Forward Logs to External Services 22
- Monitor the Firewall Using SNMP 29

Create the Security Perimeter 31

Security Perimeter Overview 32
- Firewall Deployments 32
- About Network Address Translation (NAT) 34
- About Security Policies 34

Set Up Interfaces and Zones 40
- Plan Your Deployment 40
- Configure Interfaces and Zones 41

Configure NAT Policies 43
- Translate Internal Client IP Addresses to your Public IP Address 44
- Enable Clients on the Internal Network to Access your Public Servers 45
- Enable Bi-Directional Address Translation for your Public-Facing Servers 46

Set Up Basic Security Policies 48
- Create Security Rules 48
- Test Your Security Policies 51
- Monitor the Traffic on Your Network 52

Protect Your Network Against Threats 55
Enable WildFire 56
Scan Traffic for Threats 58
Set Up Antivirus, Anti-Spyware, and Vulnerability Protection 58
Set Up File Blocking 60
Control Access to Web Content 63
Configure User Identification 67
User Identification Overview 68
About Group Mapping 68
About User Mapping 69
Enable User Identification 71
Map Users to Groups 71
Map IP Addresses to Users 73
Enable User- and Group-Based Policy 82
Verify the User-ID Configuration 84
Set Up High Availability 87
HA Overview 88
HA Modes 88
HA Links and Backup Links 88
Device Priority and Preemption 89
Failover Triggers 90
HA Timers 90
Prerequisites for Active/Passive HA 93
Configuration Guidelines 94
Configure an Active/Passive Pair 96
Define the Failover Conditions 102
Verify Failover 103

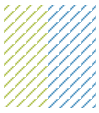

Integrate the Firewall into Your Management Network

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and prepare it for security configuration:

- Set Up Management Access to the Firewall
- Activate Firewall Services
- Add Firewall Administrators
- Monitor the Firewall

Set Up Management Access to the Firewall

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band port for managing your device going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up a data port to provide access to required external services or plan to manually upload updates regularly.

The following sections provide instructions for setting up management access to the firewall:

- Determine Your Management Strategy
- Perform Initial Configuration
- Set Up Network Access for External Services

Determine Your Management Strategy

The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using Panorama, the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:

- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage device specific configuration locally on a device and enforce shared policies across all devices or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (**ACC**) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

The procedures in this document describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, after you complete the instructions in the Perform Initial Configuration section of this guide and verify that the firewall can establish a connection to Panorama, refer to the *Panorama Administrator's Guide* for further instructions on configuring your firewall centrally.

Perform Initial Configuration

By default, the firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT interface, even if you do not plan to use this interface for your firewall management, or using a direct serial connection to the console port on the device.

Set Up Network Access to the Firewall		
Step 1	Gather the required information from your network administrator.	• IP address for MGT port • Netmask • Default gateway • DNS server address
Step 2	Connect your computer to the firewall.	You can connect to the firewall in one of the following ways: • Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the device is ready, the prompt changes to the name of the firewall, for example `PA-500 login`. • Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to `https://192.168.1.1`. Note that you may need to change the IP address on your computer to an address in the 192.168.1.0 network, such as 192.168.1.2, in order to access this URL.
Step 3	When prompted, log in to the firewall.	You must log in using the default username and password (admin/admin). The firewall will begin to initialize.
Step 4	Configure the MGT interface.	1. Select **Device > Setup > Management** and then click the Edit icon in the Management Interface Settings section of the screen. Enter the **IP Address**, **Netmask**, and **Default Gateway**. 2. Set the **Speed** to **auto-negotiate**. 3. Select which management services to allow on the interface. Best Practice: Make sure **Telnet** and **HTTP** are not selected because these services use plaintext and are not as secure as the other services. 4. Click **OK**.
Step 5	(Optional) Configure general firewall settings.	1. Select **Device > Setup > Management** and click the Edit icon in the General Settings section of the screen. 2. Enter a **Hostname** for the firewall and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain. 3. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map. 4. Click **OK**.

Set Up Network Access to the Firewall (Continued)	
Step 6 Configure DNS, time and date settings. **Note** You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP.	1. Select **Device > Setup > Services** and click the Edit icon in the Services section of the screen. 2. Enter the IP address of your **Primary DNS Server** and optionally your **Secondary DNS Server**. 3. To use the virtual cluster of time servers on the Internet, enter the hostname **pool.ntp.org** as the **Primary NTP Server** or add the IP address of your **Primary NTP Server** and optionally your **Secondary NTP Server**. 4. Click **OK** to save your settings.
Step 7 Set a secure password for the admin account.	1. Select **Device > Administrators.** 2. Select the **admin** role. 3. Enter the current default password and the new password. 4. Click **OK** to save your settings.
Step 8 Commit your changes. **Note** When the configuration changes are saved, you will lose connectivity to the web interface because the IP address will have changed.	Click **Commit**. The device may take up to 90 seconds to save your changes.
Step 9 Connect the firewall to your network.	1. Disconnect the firewall from your computer. 2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable the firewall to is configured for auto-negotiation.
Step 10 Open an SSH management session to the firewall.	Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it.

Step 11 Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server, in one of the following ways:
• If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to Set Up Network Access for External Services.
• If you do plan to allow external network access to the MGT interface, verify that you have connectivity and then proceed to Activate Firewall Services.
	If you cabled your MGT port for external network access, verify that you have access to and from the firewall by using the ping utility from the CLI. Make sure you have connectivity to the default gateway, DNS server, and the Palo Alto Networks Update Server as shown in the following example:
	`admin@PA-200> ping host updates.paloaltonetworks.com`
	`PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data.`
	`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=40.5 ms`
	`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=53.6 ms`
	`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=79.5 ms`
	Note After you have verified connectivity, press Ctrl+C to stop the pings.

Set Up Network Access for External Services

By default, the firewall uses the MGT interface to access remote services, such as DNS servers, content updates, and license retrieval. If you do not want to enable external network access to your management network, you must set up a data port to provide access to these required external services.

This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Create the Security Perimeter.

Set Up a Data Port for Access to External Services		
Step 1	Decide which port you want to use for access to external services and connect it to your switch or router port.	The interface you use will need to have a static IP address.
Step 2	Log in to the web interface.	Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page.
Step 3	(Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define.	You must delete the configuration in the following order: 1. To delete the default security policy, select **Policies > Security**, select the rule, and click **Delete**. 2. Next, delete the default virtual wire by selecting **Network > Virtual Wires**, selecting the virtual wire and clicking **Delete**. 3. To delete the default trust and untrust zones, select **Network > Zones**, select each zone and click **Delete**. 4. Finally, delete the interface configurations by selecting **Network > Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**. 5. **Commit** the changes.

Set Up a Data Port for Access to External Services (Continued)	
Step 4 Configure the interface.	1. Select **Network > Interfaces** and select the interface that corresponds to the port you cabled in Step 1. 2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**. 3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. 4. In the Zone dialog, define a **Name** for new zone, for example L3-trust, and then click **OK**. 5. Select the **IPv4** tab, select the **Static** radio button, and click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.254/24. 6. Select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. 7. Enter a **Name** for the profile, such as allow_ping, and then select the services you want to allow on the interface. These services provide management access to the device, so only select the services that correspond to the management activities you want to allow on this interface. For example, if you plan to use the MGT interface for device configuration tasks through the web interface or CLI, you would not want to enable HTTP, HTTPS, SSH, or Telnet so that you could prevent unauthorized access through this interface. For the purposes of allowing access to the external services you probably only need to enable **Ping** and then click **OK**. 8. To save the interface configuration, click **OK**.

Set Up a Data Port for Access to External Services (Continued)

Step 5 Because the firewall uses the MGT interface by default to access the external services it requires, you must change the interface the firewall uses to send these requests by editing the service routes.

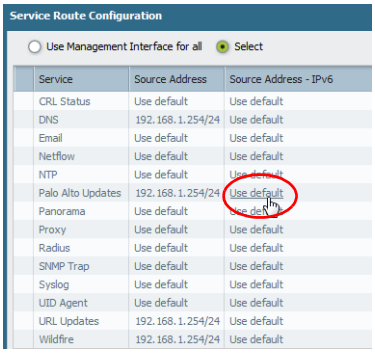

1. Select **Device > Setup > Services > Service Route Configuration**.

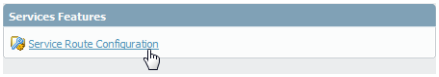

Note For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for **DNS**, **Palo Alto Updates**, **URL Updates**, and **WildFire**.

2. Click the **Customize** radio button, and select one of the following:
   - For a predefined service, select **IPv4** or **IPv6** and click the link for the service for which you want to modify the **Source Interface** and select the interface you just configured.

     If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
   - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** name and select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
3. Click **OK** to save the settings.
4. Repeat steps 2-3 above for each service route you want to modify.
5. **Commit** your changes.

Set Up a Data Port for Access to External Services (Continued)

Step 6 Configure an external-facing interface and an associated zone and then create security and NAT policy rules to allow the firewall to send service requests from the internal zone to the external zone:

1. Select **Network > Interfaces** and then select your external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as l3-untrust. You do not need to set up management services on this interface.
2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server and external DNS servers, select **Policies > Security** and click **Add**. For the purposes of initial configuration, you can create a simple rule that allows all traffic from l3-trust to l3-untrust as follows:

		Source				Destination				
Name	Tag	Zone	Address	User	HIP Profile	Zone	Address	Application	Service	Action
rule1	none	l3-trust	any	any	any	l3-untrust	any	any	any	

3. If you are using a private IP address on the internal-facing interface, you will need to create a source NAT rule to translate the address to a publicly routable address. Select **Policies > NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, l3-trust to l3-untrust in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**. For more information on NAT, see Configure NAT Policies.

		Original Packet					
Name	Tag	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service
source NAT	none	l3-trust	l3-untrust	any	any	any	any

4. **Commit** your changes.

Step 7 Verify that you have connectivity from the data port to the external services, including the default gateway, DNS server, and the Palo Alto Networks Update Server.

After you verify you have the required network connectivity, continue to Activate Firewall Services.

Launch the CLI and use the ping utility to verify that you have connectivity. Keep in mind that by default pings are sent from the MGT interface, so in this case you must specify the source interface for the ping requests as follows:

```
admin@PA-200> ping source 192.168.1.254 host
updates.paloaltonetworks.com
PING updates.paloaltonetworks.com (67.192.236.252) from
192.168.1.254 : 56(84) bytes of data.
64 bytes from 67.192.236.252: icmp_seq=1 ttl=242 time=56.7 ms
64 bytes from 67.192.236.252: icmp_seq=2 ttl=242 time=47.7 ms
64 bytes from 67.192.236.252: icmp_seq=3 ttl=242 time=47.6 ms
^C
```

After you have verified connectivity, press Ctrl+C to stop the pings.

Activate Firewall Services

Before you can begin using the firewall to secure your network, you must register it and activate the licenses for the services you have purchased. In addition, you should ensure that you are running the appropriate version of PAN-OS as described in the following sections:

- Register With Palo Alto Networks
- Activate Licenses
- Manage Content Updates
- Install Software Updates

Register With Palo Alto Networks

Register the Firewall		
Step 1	Log in to the web interface.	Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page.
Step 2	Locate your serial number and copy it to the clipboard.	On the **Dashboard**, locate your **Serial Number** in the General Information section of the screen.
Step 3	Go to the Palo Alto Networks Support site.	In a new browser tab or window, go to https://support.paloaltonetworks.com.
Step 4	Register the device. The way you register depends on whether you already have a login to the support site. 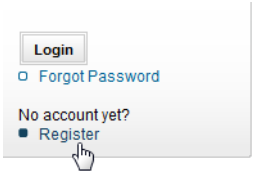	• If this is the first Palo Alto Networks device you are registering and you do not yet have a login, click **Register** on the right side of the page. To register, you must provide your sales order number or customer ID, and the serial number of your firewall (which you can paste from your clipboard) or the authorization code you received with your order. You will also be prompted to set up a username and password for access to the Palo Alto Networks support community. • If you already have a support account, log in and then click **My Devices**. Scroll down to Register Device section at the bottom of the screen and enter the serial number of your firewall (which you can paste from your clipboard), your city and postal code and then click **Register Device**.

Activate Licenses

Before you can start using your firewall to secure the traffic on your network, you must activate the licenses for each of the services you purchased. Available licenses and subscriptions include the following:

- **Threat Prevention**—Provides antivirus, anti-spyware, and vulnerability protection. For more information about threat prevention, see Set Up Antivirus, Anti-Spyware, and Vulnerability Protection.
- **Decryption Port Mirror**—Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures-such as NetWitness or Solera-for archiving and analysis.
- **URL Filtering**—In order to create policy rules based on dynamic URL categories, you must purchase and install a subscription for one of the supported URL filtering databases: PAN-DB or BrightCloud. For more information about URL filtering, see Control Access to Web Content.
- **Virtual Systems**—This license is required to enable support for multiple virtual systems on PA-2000 and PA-3000 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (the base number varies by platform). The PA-500, PA-200, and VM-Series firewalls do not support virtual systems.
- **WildFire**—Although basic WildFire support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, enabling sub-hourly WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to a private WF-500 WildFire appliance. For more information about WildFire, see Enable WildFire.
- **GlobalProtect**—Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy a single GlobalProtect portal and gateway (without HIP checks) without a license. However, if you want to deploy multiple gateways, you must purchase a portal license (one-time, permanent license). If you want to use host checks you will also need gateway licenses (subscription) for each gateway. For more information on GlobalProtect, refer to the GlobalProtect Administrator's Guide.

Activate Licenses		
Step 1	Locate the activation codes for the licenses you purchased.	When you purchased your subscriptions you should have received an email from Palo Alto Networks customer service listing the activation code associated with each subscription. If you cannot locate this email, contact customer support to obtain your activation codes before you proceed.
Step 2	Launch the web interface and go to the license page.	Select **Device > Licenses**.
Step 3	Activate each license you purchased.	1. Select **Activate feature using authorization code**. 2. When prompted, enter the **Authorization Code** and then click **OK**. 3. Verify that the license was successfully activated. For example, after activating the WildFire license, you should see that the license is valid:
Note	If your firewall does not have Internet access from the management port, you can manually download your license files from the support site and upload them to your firewall using the **Manually upload license key** option.	

Manage Content Updates

In order to stay ahead of the changing threat and application landscape, all Palo Alto Networks firewalls support dynamic content updates. Depending on which subscriptions you've purchased, these updates include the latest application and threat signatures, along with a URL filtering database. To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure that you keep your firewalls up-to-date with the latest updates published by Palo Alto Networks. The following content updates are available, depending on which subscriptions you have:

Although you can manually download and install content updates at any time, as a best practice you should schedule updates to occur automatically.

- **Antivirus**—Includes new and updated antivirus signatures, including signatures discovered by the WildFire cloud service. You must have a Threat Prevention subscription to get these updates. New antivirus signatures are published daily.
- **Applications**—Includes new and updated application signatures. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published weekly.
- **Applications and Threats**—Includes new and updated application and threat signatures. This update is available if you have a Threat Prevention subscription (and you get it instead of the Applications update). New Applications and Threats updates are published weekly.
- **GlobalProtect Data File**—Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect agents. You must have a GlobalProtect portal and GlobalProtect gateway license in order to receive these updates. In addition, you must create a schedule for these updates before GlobalProtect will function.
- **BrightCloud URL Filtering**—Provides updates to the BrightCloud URL Filtering database only. You must have a BrightCloud subscription to get these updates. New BrightCloud URL database updates are published daily. If you have a PAN-DB license, scheduled updates are not required as devices remain in-sync with the servers automatically.
- **WildFire**—Provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire cloud service. Without the subscription, you must wait 24 to 48 hours for the signatures to roll into the Applications and Threat update.

If your firewall does not have Internet access from the management port, you can download content updates from the Palo Alto Networks Support Site (https://support.paloaltonetworks.com) and then Upload them to your firewall.

Download the Latest Databases		
Step 1	Launch the web interface and go to the Dynamic Updates page.	Select **Device > Dynamic Updates**.

Download the Latest Databases (Continued)

Step 2 Check for the latest updates.

Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available:

- **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**.

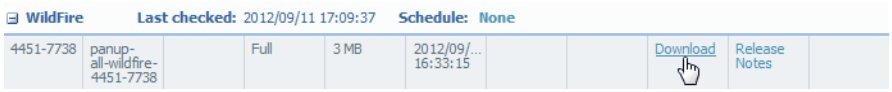

Note You cannot download the antivirus database until you have installed the Application and Threats database.

- **Upgrade**—Indicates that there is a new version of the BrightCloud database available. Click the link to begin the download and installation of the database. The database upgrade begins in the background; when completed a check mark displays in the **Currently Installed** column. Note that if you are using PAN-DB as your URL filtering database you will not see an upgrade link because the PAN-DB database automatically stays in sync with the server.

Tip: To check the status of an action, click **Tasks** (on the lower right-hand corner of the window).

- **Revert**—Indicates that the corresponding software version has been downloaded previously. You can choose to revert to the previously installed version of the update.

Step 3 Install the updates.

Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column.

Note Installation can take up to 20 minutes on a PA-200, PA-500, or PA-2000 device and up to two minutes on a PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, or VM-Series firewall.

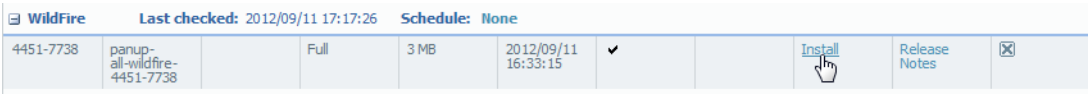

Download the Latest Databases (Continued)		
Step 4	Schedule each update. Repeat this step for each update you want to schedule. Best Practice: Stagger the update schedules because the firewall can only download one update at a time. If you schedule the updates to download during the same time interval, only the first download will succeed.	1. Set the schedule of each update type by clicking the **None** link. 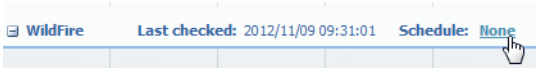 2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The available values vary by content type (WildFire updates are available **Every 15 minutes**, **Every 30 minutes** or **Every Hour** whereas all other content types can be scheduled for **Daily** or **Weekly** update). 3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur. 4. Specify whether you want the system to **Download And Install** the update (best practice) or **Download Only**. 5. In rare instances, errors in content updates may be found. For this reason, you may want to delay installing new updates until they have been released for a certain number of hours. You can specify how long after a release to wait before performing a content update by entering the number of hours to wait in the **Threshold (Hours)** field. 6. Click **OK** to save the schedule settings. 7. Click **Commit** to save the settings to the running configuration.

Install Software Updates

When installing a new firewall, it is a good idea to upgrade to the latest software update (or to the update version recommended by your reseller or Palo Alto Networks Systems Engineer) to take advantage of the latest fixes and security enhancements. Note that before updating the software, you should first make sure you have the latest content updates as detailed in the previous section (the release notes for a software update specify the minimum content update versions that are supported in the release).

Update PAN-OS		
Step 1	Launch the web interface and go to the Software page.	Select **Device > Software**.
Step 2	Check for software updates.	Click **Check Now** to check for the latest updates. If the value in the **Action** column is **Download** it indicates that an update is available.

Update PAN-OS (Continued)

Step 3 Download the update.

Note If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Site (https://support.paloaltonetworks.com). You can then manually **Upload** them to your firewall.

Locate the version you want and then click **Download**. When the download completes, the value in the Action column changes to **Install**.

Version	Size	Release Date	Downloaded	Currently Installed	Action
5.0.0	259 MB	2012/11/01 19:58:24	✔		Install
4.1.9	169 MB	2012/11/05 23:40:31			Download
4.1.8	168 MB	2012/09/22 21:01:08	✔	✔	Download
4.1.8-h3	168 MB	2012/10/18 23:49:21			Download
4.1.7	152 MB	2012/07/29			Download

Step 4 Install the update.

1. Click **Install**.
2. Reboot the firewall:
   - If you are prompted to reboot, click **Yes**.

   - If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section of the screen.

Add Firewall Administrators

By default, every Palo Alto Networks firewall comes preconfigured with a default administrative account (admin), which provides full read-write access (also known as superuser access) to the firewall.

As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the firewall from unauthorized configuration (or modification) and to enable logging of the actions of each individual firewall administrator.

The following sections describe the various ways you can set up administrative accounts and provide procedures for setting up basic administrative access:

- Administrative Roles
- Administrative Authentication
- Create an Administrative Account

Administrative Roles

The way you configure administrator accounts depends on the security requirements within your organization, whether you have existing authentication services you want to integrate with, and how many different administrative roles you require. A *role* defines the type of access the associated administrator has to the system. There are two types of roles you can assign:

- **Dynamic Roles**—Built-in roles that provide Superuser, Superuser (read-only), Device administrator, Device administrator (read-only), Virtual system administrator, and Virtual system administrator (read-only) access to the firewall. With dynamic roles, you don't have to worry about updating the role definitions as new features are added because the roles automatically update.
- **Admin Role Profiles**—Allow you to create your own role definitions in order to provide more granular access control to the various functional areas of the web interface, CLI and/or XML API. For example, you could create an Admin Role Profile for your operations staff that provides access to the device and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definition, logs, and reports. Keep in mind that with Admin Role Profiles you must update the profiles to explicitly assign privileges for new features/components that are added to the product.

Administrative Authentication

There are four ways you can authenticate administrative users:

- **Local administrator account with local authentication**—Both the administrator account credentials and the authentication mechanisms are local to the firewall. You can further secure the local administrator account by creating a password profile that defines a validity period for passwords and by setting device-wide password complexity settings.

- **Local administrator account with SSL-based authentication**—With this option, you create the administrator accounts on the firewall, but authentication is based on SSH certificates (for CLI access) or client certificates/common access cards (for the web interface). Refer to the article How to Configure Certificate-based Authentication for the WebUI for details on how to configure this type of administrative access.
- **Local administrator account with external authentication**—The administrator accounts are managed on the local firewall, but the authentication functions are offloaded to an existing LDAP, Kerberos, or RADIUS service. To configure this type of account, you must first create an authentication profile that defines how to access the external authentication service and then create an account for each administrator that references the profile.
- **External administrator account and authentication**—Account administration and authentication are handled by an external RADIUS server. To use this option, you must define Vendor Specific Attributes (VSAs) on your RADIUS server that map to the admin role and, optionally, the virtual system objects you have defined on the Palo Alto Networks device. Refer to the Radius Vendor Specific Attributes (VSA) article for details on how to configure this type of administrative access.

Create an Administrative Account

Create administrative accounts to define access and administrative privileges for firewall administrators. Because it is common to delegate specific administrative tasks to specific administrators with varying roles, Palo Alto Networks recommends that you create admin role profiles that allow administrators access only to the areas of the management interface that are required to perform their jobs. You can assign the various roles you create to individual administrator accounts and specify access privileges to each management interface: the web interface, the Command Line Interface (CLI), and the REST Management API. By creating admin roles with very granular access privileges, you can ensure that sensitive company data is protected and end user privacy is ensured.

The following procedure describes how to create a local administrator account with local authentication, including how to set administrator access for each management interface.

Create a Local Administrator	
Step 1 If you plan to use Admin Role Profiles rather than Dynamic Roles, create the profiles that define what type of access, if any, to give to the different sections of the web interface, CLI, and XML API for each administrator assigned to the role.	Complete the following steps for each role you want to create: 1. Select **Device > Admin Roles** and then click **Add**. 2. Enter a **Name** and optionally a **Description** for the role. 3. On the Web UI, Command Line and/or XML API tabs, specify the access to allow for each management interface: • On the **Web UI** and/or **XML API** tabs, set the access levels for each functional area of the interface by clicking the icon to toggle it to the desired setting: Enable, Read Only, or Disable. • On the **Command Line** tab, specify the type of access to allow to the CLI: **superreader**, **deviceadmin**, or **devicereader** (for Device roles); **vsysadmin** or **vsysreader** (for Virtual System roles); or **None** to disable CLI access entirely. 4. Click **OK** to save the profile. For example, allow an admin full access to a device using the XML API, with the exception of importing or exporting files: 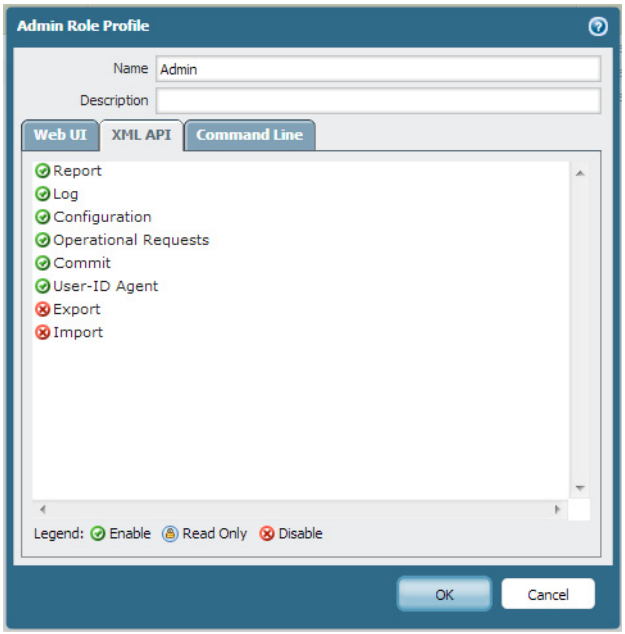

Create a Local Administrator (Continued)		
Step 2	(Optional) Set requirements for local user-defined passwords.	• **Create Password Profiles**—Define how often administrators must change their passwords. You can create multiple password profiles and apply them to administrator accounts as needed to enforce the desired security. To create a password profile, select **Device > Password Profiles** and then click the **Add**. • **Configure minimum password complexity settings**—Define rules that govern password complexity, allowing you to force administrators to create passwords that are harder to guess, crack, or compromise. Unlike password profiles, which can be applied to individual accounts, these rules are device wide and apply to all passwords. To configure the settings, select **Device > Setup** and then click the Edit icon in the Minimum Password Complexity section.
Step 3	Create an account for each administrator.	1. Select **Device > Administrators** and then click **Add**. 2. Enter a user **Name** and **Password** for the administrator. 3. Select the **Role** to assign to this administrator. You can either select one of the predefined Dynamic roles or a custom Role Based profile if you created one in Step 1. 4. (Optional) Select a **Password Profile**. 5. Click **OK** to save the account.
Step 4	Commit your changes.	1. Click **Commit**.

Monitor the Firewall

Another thing to consider during your initial deployment is how you plan to monitor the firewall—both for proper functioning of the firewall as well as the monitoring of the traffic and threats it manages and controls. Do you have centralized services, such as Syslog or SNMP, that you want to leverage? Do you have specific log file archive, auditing, and/or backup requirements?

The following sections describe the methods you can use to monitor the firewall and provide basic setup instructions:

- Monitor Applications and Threats
- View Local Log Data
- Forward Logs to External Services
- Monitor the Firewall Using SNMP

You can also configure the firewall (excluding PA-4000 Series and PA-7050 firewalls) to export flow data to a NetFlow collector for analysis and reporting.

Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the App-ID technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then monitor the applications from the Application Command Center (**ACC**). The ACC graphically summarizes the log database to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC.

You can quickly investigate new, risky, or unfamiliar applications that appear in ACC with a single click that displays a description of the application, its key features, its behavioral characteristics, and who is using it. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy.

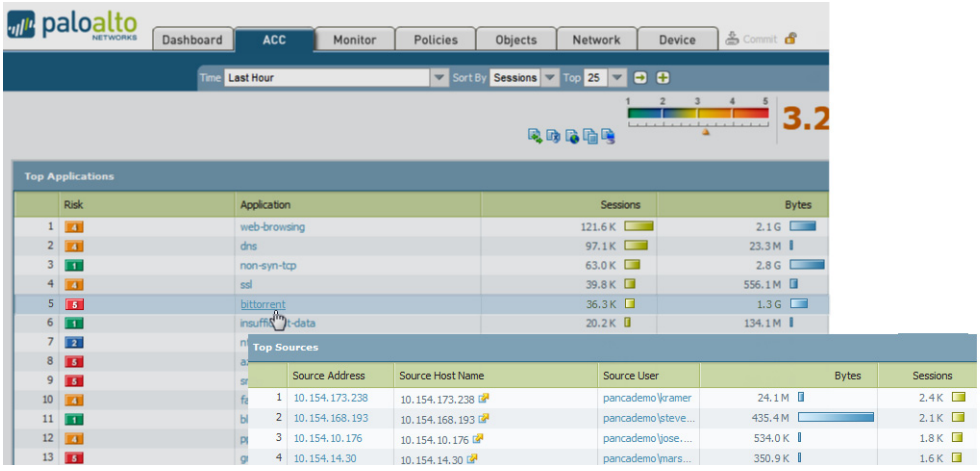

View Local Log Data

All Palo Alto Networks next-generation firewalls can generate log files that provide an audit trail of the activities and events on the firewall. There are separate logs for separate types of activities and events. For example, the Threat logs record all traffic that causes the firewall to generate a security alarm, whereas URL Filtering logs record all traffic that matches a URL Filtering profile attached to a security policy, and Config logs record all changes to the firewall configuration.

There are several ways you can view the log data on the local firewall:

- View the Log Files
- Display Log Data on the Dashboard
- View Reports

View the Log Files

By default all log files are generated and stored locally on the firewall. You can view these log files directly (**Monitor > Logs**):

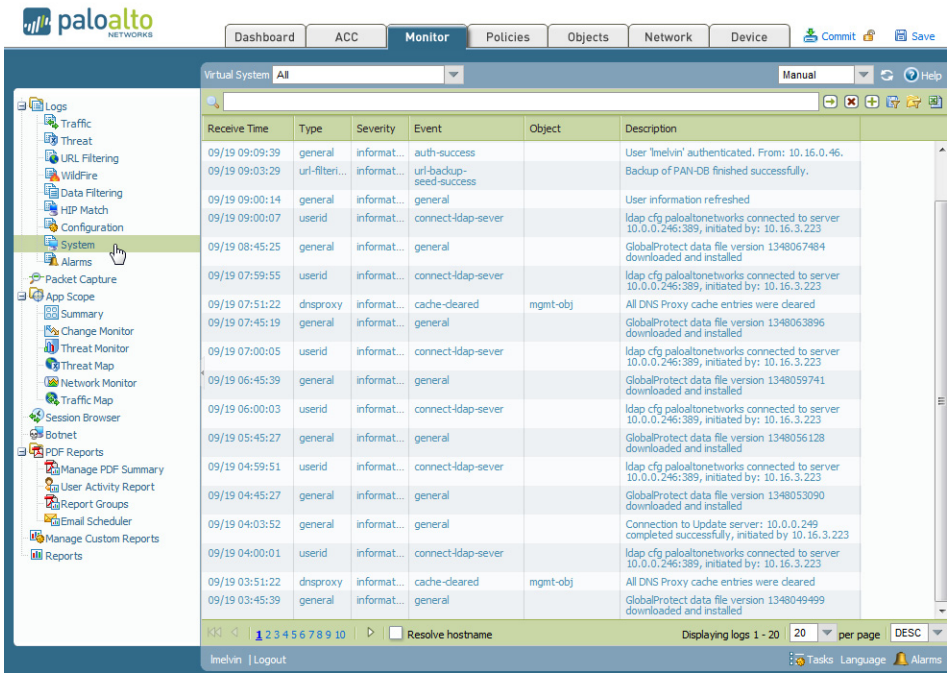
paloalto
Dashboard
ACC
Monitor
Policies
Objects
Network
Device
Commit
Save
Virtual System All
Manual
Help
Logs
Traffic
Threat
URL Filtering
WildFire
Data Filtering
HIP Match
Configuration
System
Alarms
Packet Capture
App Scope
Summary
Change Monitor
Threat Monitor
Threat Map
Network Monitor
Traffic Map
Session Browser
Botnet
PDF Reports
Manage PDF Summary
User Activity Report
Report Groups
Email Scheduler
Manage Custom Reports
Reports
Receive Time
Type
Severity
Event
Object
Description
09/19 09:09:39 general informat... auth-success User 'lmelvin' authenticated. From: 10.16.0.46.
09/19 09:03:29 url-filteri... informat... url-backup-seed-success Backup of PAN-DB finished successfully.
09/19 09:00:14 general informat... general User information refreshed
09/19 09:00:07 userid informat... connect-ldap-sever ldap cfg paloaltonetworks connected to server 10.0.0.246:389, initiated by: 10.16.3.223
09/19 08:45:25 general informat... general GlobalProtect data file version 1348067484 downloaded and installed
09/19 07:59:55 userid informat... connect-ldap-sever ldap cfg paloaltonetworks connected to server 10.0.0.246:389, initiated by: 10.16.3.223
09/19 07:51:22 dnsproxy informat... cache-cleared mgmt-obj All DNS Proxy cache entries were cleared
09/19 07:45:19 general informat... general GlobalProtect data file version 1348063896 downloaded and installed
09/19 07:00:05 userid informat... connect-ldap-sever ldap cfg paloaltonetworks connected to server 10.0.0.246:389, initiated by: 10.16.3.223
09/19 06:45:39 general informat... general GlobalProtect data file version 1348059741 downloaded and installed
09/19 06:00:03 userid informat... connect-ldap-sever ldap cfg paloaltonetworks connected to server 10.0.0.246:389, initiated by: 10.16.3.223
09/19 05:45:27 general informat... general GlobalProtect data file version 1348056128 downloaded and installed
09/19 04:59:51 userid informat... connect-ldap-sever ldap cfg paloaltonetworks connected to server 10.0.0.246:389, initiated by: 10.16.3.223
09/19 04:45:27 general informat... general GlobalProtect data file version 1348053090 downloaded and installed
09/19 04:03:52 general informat... general Connection to Update server: 10.0.0.249 completed successfully, initiated by 10.16.3.223
09/19 04:00:01 userid informat... connect-ldap-sever ldap cfg paloaltonetworks connected to server 10.0.0.246:389, initiated by: 10.16.3.223
09/19 03:51:22 dnsproxy informat... cache-cleared mgmt-obj All DNS Proxy cache entries were cleared
09/19 03:45:39 general informat... general GlobalProtect data file version 1348049499 downloaded and installed
1 2 3 4 5 6 7 8 9 10
Resolve hostname
Displaying logs 1 - 20
20
per page
DESC
lmelvin | Logout
Tasks Language Alarms

Display Log Data on the Dashboard

You can also monitor the local log data directly from the **Dashboard** by adding the associated widgets:

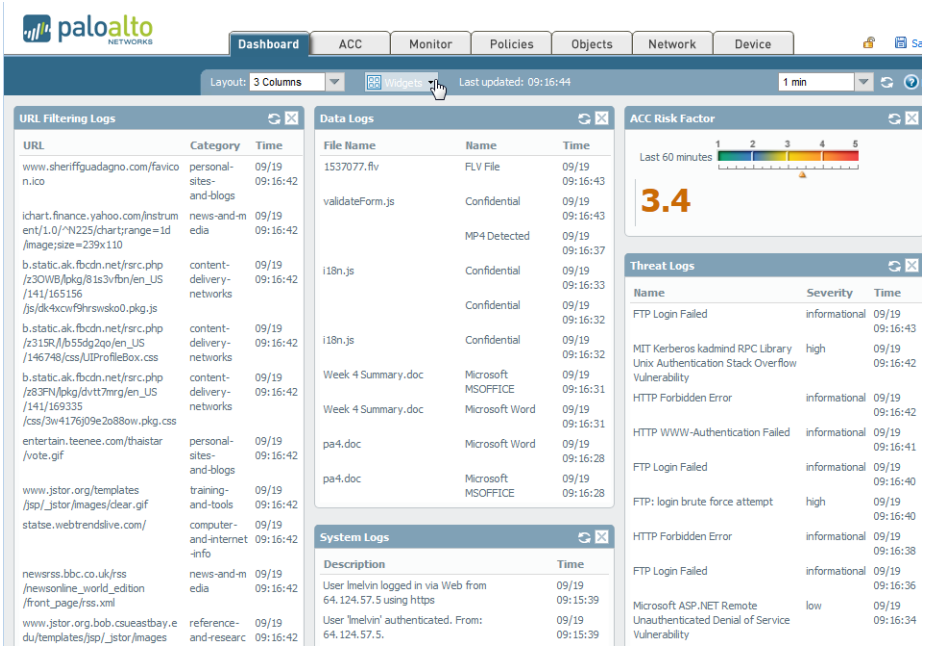

View Reports

The firewall also uses the log data to generate reports (**Monitor > Reports**) that display the log data in a tabular or graphical format.

Forward Logs to External Services

Depending on the type and severity of the data in the log files, you may want to be alerted to critical events that require your attention, or you may have policies that require you to archive the data for longer than it can be stored on the firewall. In these cases you will want to forward your log data to an external service for archive, notification, and/or analysis.

When configuring log forwarding or WildFire file forwarding on a PA-7050 firewall, you must configure one data port with the log type **Log Card**. This port must be connected to a network that can communicate with your logging destinations and if WildFire forwarding is configured, it must be able to reach the WildFire cloud or a WildFire appliance if one is deployed on your network. After configuring an interface with this type, these services will automatically use the log card port and no other configuration is required. For more information on this interface type, see the help in the Network > Interfaces section of a PA-7050.

To forward log data to an external service you must complete the following tasks:

- Configure the firewall to access the remote services that will be receiving the logs. See Define Remote Logging Destinations.
- Configure each log type for forwarding. See Enable Log Forwarding.

Define Remote Logging Destinations

In order to reach an external service—such as a Syslog server or SNMP trap manager—the firewall must know the details of how to access and, if necessary, authenticate to the service. On the firewall, you define this information in a Server Profile. You must create a Server Profile for each external service you want the firewall to interact with. The type of logging destinations you need to set up and which logs you forward will depend on your needs. Some common log forwarding scenarios include the following:

- For immediate notification about critical system events or threats that require your attention, you can generate SNMP traps or send email alerts. See Set Up Email Alerts and/or Set Up SNMP Trap Destinations.
- For long-term storage and archival of data and for centralized device monitoring, you can send the log data to a Syslog server. See Define Syslog Servers. This enables integration with third-party security monitoring tools, such as Splunk! or ArcSight.

If you do not have a Syslog collector or if you do not require real-time updates, you can instead schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the firewall and a remote host. For more information, refer to the "Reports and Logging" section of the *PAN-OS 6.0 Administrator's Guide.*

- For aggregation and reporting of log data from multiple Palo Alto Networks firewalls, you can forward logs to a Panorama Manager or Panorama Log Collector. See Forward Logs to Panorama.

You can define as many Server Profiles as you need. For example, you could use separate Server Profiles to send traffic logs to one Syslog server and system logs to a different one. Or, you could include multiple server entries in a single Server Profile to enable you to log to multiple Syslog servers for redundancy.

By default, all log data is forwarded over the MGT interface. If you plan to use an interface other than MGT, you will need to configure a Service Route for each service to which you plan to forward logs as described in Step 5 of the procedure to Set Up a Data Port for Access to External Services.

Set Up Email Alerts

Set Up Email Alerts		
Step 1	Create a Server Profile for your email server.	1. Select **Device > Server Profiles > Email**. 2. Click **Add** and then enter a **Name** for the profile. 3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down. 4. Click **Add** to add a new email server entry and enter the information required to connect to the Simple Mail Transport Protocol (SMTP) server and send email (you can add up to four email servers to the profile): • **Name**—Name to identify the mail server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server. • **Email Display Name**—The name to show in the From field of the email. • **From**—The email address where notification emails will be sent from. • **To**—The email address to which notification emails will be sent. • **Additional Recipient**—If you want the notifications sent to a second account, enter the additional address here. You can only add one additional recipient. To add multiple recipients, add the email address of a distribution list. • **Email Gateway**—The IP address or host name of the SMTP gateway to use to send the emails. 5. Click **OK** to save the server profile.
Step 2	(Optional) Customize the format of the email messages the firewall sends.	Select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.
Step 3	Save the server profile and commit your changes.	1. Click **OK** to save the profile. 2. Click **Commit** to save the changes to the running configuration.

Set Up SNMP Trap Destinations

Simple Network Management Protocol (SNMP) is a standard facility for monitoring the devices on your network. You can configure the firewall to send SNMP traps to your SNMP management software to alert you to critical system events or threats that require your immediate attention.

You can also use SNMP to monitor the firewall. In this case, your SNMP manager must be configured to get statistics from the firewall rather than (or in addition to) having the firewall send traps to the manager. For more information, see Monitor the Firewall Using SNMP.

Set Up SNMP Trap Destinations		
Step 1	(SNMP v3 only) Get the engine ID for the firewall.	In order to find out the firewall's engine ID, you must configure the firewall for SNMP v3 and send a GET message from your SNMP manager or MIB browser as follows: 1. Enable the interface to allow inbound SNMP requests: • If you will be receiving SNMP GET messages on the MGT interface, select **Device > Setup > Management** and click the Edit icon in the Management Interface Settings section of the screen. In the **Services** section, select the **SNMP** check box and then click **OK**. • If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management. 2. Configure the firewall for SNMP v3 as described in Step 2 in Set Up SNMP Monitoring. If you do not configure the firewall for SNMP v3 your MIB browser will not allow you to GET the engine ID. 3. Connect your MIB browser or SNMP manager to the firewall and run a GET for OID 1.3.6.1.6.3.10.2.1.1.0. The value that is returned is the unique engine ID for the firewall.
Note	In many cases, the MIB browser or SNMP manager will automatically discover the engine ID upon successful connection to the SNMP agent on the firewall. You can usually find this information in the agent settings section of the interface. Refer to the documentation for your specific product for instructions on finding the agent information.	

Set Up SNMP Trap Destinations (Continued)

Step 2 Create a Server Profile that contains the information for connecting and authenticating to the SNMP manager(s).

1. Select **Device > Server Profiles > SNMP Trap**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
4. Specify the version of SNMP you are using (V2c or V3).
5. Click **Add** to add a new **SNMP Trap Receiver** entry (you can add up to four trap receivers per server profile). The required values depend on whether you are using SNMP V2c or V3 as follows:

   SNMP V2c

   - **Name**—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server.
   - **SNMP Manager**—The IP address of the SNMP manager to which you want to send traps.
   - **Community**—The community string required to authenticate to the SNMP manager.

   SNMP V3

   - **Name**—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server.
   - **SNMP Manager**—The IP address of the SNMP manager to which you want to sent traps.
   - **User**—The username required to authenticate to the SNMP manager.
   - **EngineID**—The engine ID of the firewall, as identified in Step 1. This is a hexadecimal value from 5 to 64 bytes with a 0x prefix. Each firewall has a unique engine ID.
   - **Auth Password**—The password to be used for authNoPriv level messages to the SNMP manager. This password will be hashed using Secure Hash Algorithm (SHA-1), but will not be encrypted.
   - **Priv Password**—The password to be used for authPriv level messages to the SNMP manager. This password be hashed using SHA and will be encrypted using Advanced Encryption Standard (AES 128).
6. Click **OK** to save the server profile.

Step 3	(Optional) Set up a service route for SNMP traps.	By default, SNMP traps are sent over the MGT interface. If you want to use a different interface for SNMP traps, you must edit the service route to enable the firewall to reach your SNMP manager. See Set Up Network Access for External Services for instructions.
Step 4	Commit your changes.	Click **Commit**. The device may take up to 90 seconds to save your changes.
Step 5	Enable the SNMP manager to interpret the traps it receives from the firewall.	Load the PAN-OS MIB files into your SNMP management software and compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this.

Define Syslog Servers

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archive, analysis, and reporting.

There are five log types that PAN-OS can export to a Syslog server: traffic, threat, HIP match, config, and system. For more details about the fields in each log type, refer to the PAN-OS Syslog Integration Tech Note. For a partial list of log messages and their severity levels, refer to the System Log Reference.

Syslog messages are sent in clear text and cannot be directly encrypted. However, if you need encryption, you can send the Syslog messages through a tunnel interface, which will force the Syslog packets to be encrypted. You will also need to create a new service route for Syslog.

Set Up Syslog Forwarding		
Step 1	Create a Server Profile that contains the information for connecting to the Syslog server(s).	1. Select **Device > Server Profiles > Syslog**. 2. Click **Add** and then enter a **Name** for the profile. 3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down. 4. Click **Add** to add a new Syslog server entry and enter the information required to connect to the Syslog server (you can add up to four Syslog servers to the same profile): • **Name**—Unique name for the server profile. • **Syslog Server**—IP address or fully qualified domain name (FQDN) of the Syslog server. • **Transport**—Select TCP, UDP, or SSL as the method of communication with the syslog server. • **Port**—The port number on which to send Syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the Syslog server. • **Format**—Select the Syslog message format to use, BSD or IETF. Traditionally, BSD format is over UDP and IETF format is over TCP/SSL. • **Facility**—Select one of the Syslog standard values, which is used to calculate the priority (PRI) field in your Syslog server implementation. You should select the value that maps to how you use the PRI field to manage your Syslog messages. 5. (Optional) To customize the format of the Syslog messages the firewall sends, select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide. 6. Click **OK** to save the server profile.

Set Up Syslog Forwarding (Continued)	
Step 2 (Optional) Configure the hostname format in the Syslog header messages. 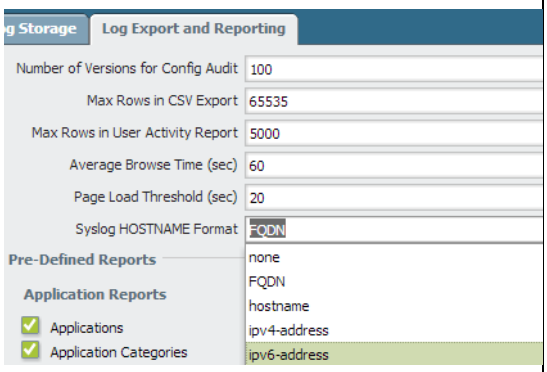	1. Select **Device > Setup > Management** and click the Edit icon in the **Logging and Reporting Settings** section. 2. Select **Log Export and Reporting** and in **Syslog HOSTNAME format**, select whether to include the hostname, FQDN (the default that concatenates the hostname and domain name) or the IPv4/IPv6 address of the device. 3. Click **OK**.
Step 3 Commit your changes.	Click **Commit**. The device may take up to 90 seconds to save your changes.

Forward Logs to Panorama

Before you can forward log files to a Panorama Manager or a Panorama Log Collector, the firewall must be configured as a managed device. For details on setting up Panorama and adding devices, refer to the *Panorama Administrator's Guide*. You can then enable log forwarding to Panorama for each type of log as described in Enable Log Forwarding. For logs forwarded to Panorama, support for centralized log forwarding to an external syslog server is available.

Enable Log Forwarding

After you create the Server Profiles that define where to send your logs, you must enable log forwarding. For each log type, you can specify whether to forward it to Syslog, email, SNMP trap receiver, and/or Panorama. The way you enable forwarding depends on the log type:

- **Traffic Logs**—Enable forwarding of Traffic logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) and adding it to the security policies you want to trigger the log forwarding. Only traffic that matches a specific rule within the security policy will be logged and forwarded.
- **Threat Logs**—Enable forwarding of Threat logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) that specifies which severity levels you want to forward and then adding it to the security policies for which you want to trigger the log forwarding. A Threat log entry will only be created (and therefore forwarded) if the associated traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection). The following table summarizes the threat severity levels:

Severity	Description
Critical	Serious threats such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions.
High	Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool.
Medium	Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access.
Low	Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. Data Filtering profile matches are logged as Low.
Informational	Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. URL Filtering log entries with a benign verdict are logged as Informational.

- **Config Logs**—Enable forwarding of Config logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > Config Logs**).
- **System Logs**—Enable forwarding of System logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > System Logs**). You must select a Server Profile for each severity level you want to forward. For a partial list of system log messages and their corresponding severity levels, refer to the System Log Reference. The following table summarizes the system log severity levels:

Severity	Description
Critical	Hardware failures, including HA failover and link failures.
High	Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers.
Medium	Mid-level notifications, such as antivirus package upgrades.
Low	Minor severity notifications, such as user password changes.
Informational	Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels.

- **WildFire Logs**—Enable forwarding of WildFire logs that contain information about files that are forwarded from the firewall to WildFire for analysis. You can configure the firewall to forward logs for the verdict Benign and/or Malicious.

Monitor the Firewall Using SNMP

All Palo Alto Networks firewalls support standard networking SNMP management information base (MIB) modules as well as proprietary Enterprise MIB modules. You can configure an SNMP manager to get statistics from the firewall. For example, you could configure your SNMP manager to monitor the interfaces, active sessions, concurrent sessions, session utilization percentage, temperature, and/or system uptime on the firewall.

Palo Alto Networks firewalls support SNMP GET requests only; SNMP SET requests are not supported.

Set Up SNMP Monitoring		
Step 1	Enable the interface to allow inbound SNMP requests.	• If you will be receiving SNMP GET messages on the MGT interface, select **Device > Setup > Management** and click the Edit icon in the Management Interface Settings section of the screen. In the **Services** section, select the **SNMP** check box and then click **OK**. • If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management.
Step 2	From the web interface on the firewall, configure the settings to allow the SNMP agent on the firewall to respond to incoming GET requests from the SNMP manager.	1. Select **Device > Setup > Operations > SNMP Setup**. 2. Specify the **Physical Location** of the firewall and the name or email address of an administrative **Contact**. 3. Select the SNMP **Version** and then enter the configuration details as follows (depending on which SNMP version you are using) and then click **OK**: • **V2c**—Enter the **SNMP Community String** that will allow the SNMP manager access to the SNMP agent on the firewall. The default value is **public**, however because this is a well-known community string, it is a best practice to use a value that is not easily guessed. • **V3**—You must create at least one View and one User in order to use SNMPv3. The view specifies which management information the manager has access to. If you want to allow access to all management information, just enter the top-level **OID** of .1.3.6.1 and specify the **Option** as **include** (you can also create views that exclude certain objects). Use **0xf0** as the **Mask**. Then when you create a user, select the **View** you just created and specify the **Auth Password** and **Priv Password**. The authentication settings (the community string for V2c or the username and passwords for V3) configured on the firewall must match the value configured on the SNMP manager. 4. Click **OK** to save the settings. 5. Click **Commit** to save the SNMP settings.

Set Up SNMP Monitoring (Continued)		
Step 3	Enable the SNMP manager to interpret firewall statistics.	Load the PAN-OS MIB files into your SNMP management software and, if necessary, compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this.
Step 4	Identify the statistics you want to monitor.	Using a MIB browser, walk the PAN-OS MIB files to identify the object identifiers (OIDs) that correspond to the statistics you want to monitor. For example, suppose you want to monitor Session Utilization Percentage on the firewall. Using a MIB browser you will see that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in the PAN-COMMON-MIB.
Step 5	Configure the SNMP management software to monitor the OIDs you are interested in.	Refer to the documentation for your SNMP manager for specific instructions on how to do this.
Step 6	After you complete the configuration on both the firewall and the SNMP manager, you can begin monitoring the firewall from your SNMP management software.	The following is an example of how an SNMP manager displays real-time session utilization percentage statistics for a monitored PA-500 firewall: 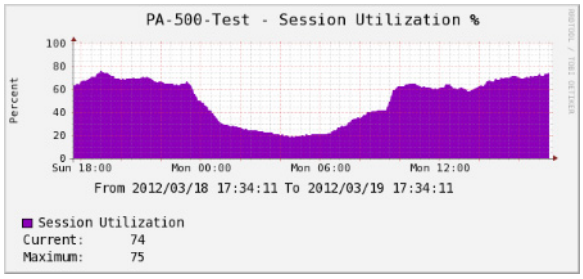

Create the Security Perimeter

The following topics provide basic steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy:

- Security Perimeter Overview
- Set Up Interfaces and Zones
- Configure NAT Policies
- Set Up Basic Security Policies

Security Perimeter Overview

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall decides how to act on a packet based on whether the packet matches a *security policy*. At the most basic level, the security policy must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, security policies are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that provides an abstraction for an area of trust for simplified policy enforcement. For example, in the following topology diagram, there are three zones: Trust, Untrust, and the Demilitarized Zone (DMZ). Traffic can flow freely within a zone, but traffic will not be able to flow between zones until you define a security policy that allows it.

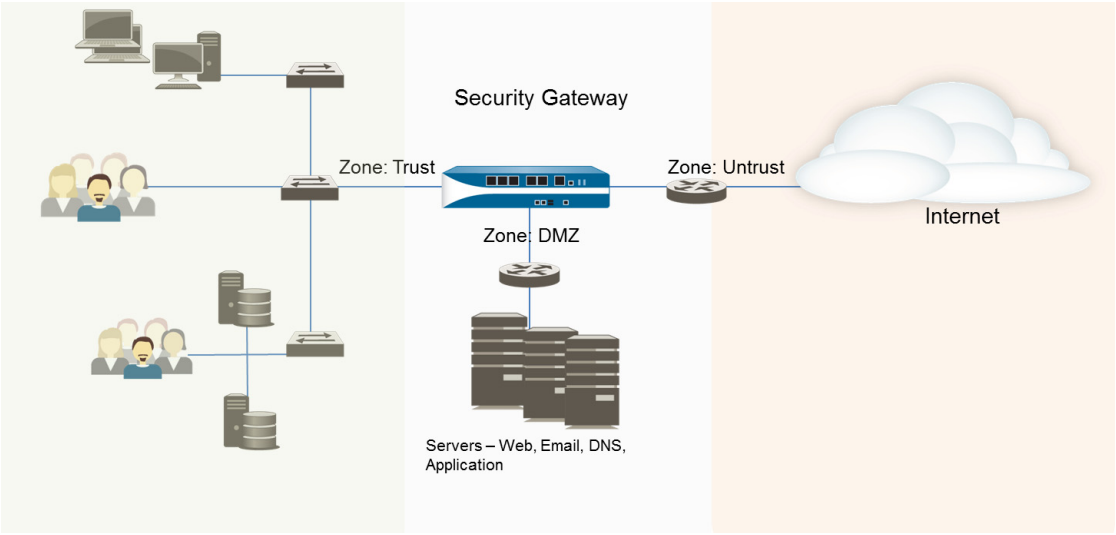

The following sections describe the components of the security perimeter and provide steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy that allows traffic from your internal zone to the Internet and to the DMZ. By initially creating a basic policy like this, you will be able to analyze the traffic running through your network and use this information to define more granular policies for safely enabling applications while preventing threats.

- Firewall Deployments
- About Network Address Translation (NAT)
- About Security Policies

Firewall Deployments

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The following sections provide basic information on each type of deployment. For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls.

Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together. By using a virtual wire, you can install the firewall in any network environment without reconfiguring adjacent devices. If necessary, a virtual wire can block or allow traffic based on the virtual LAN (VLAN) tag values. You can also create multiple subinterfaces and classify traffic according to an IP Address (address, range, or subnet), VLAN, or a combination of the two.

By default, the virtual wire "default-vwire" binds together Ethernet ports 1 and 2 and allows all untagged traffic. Choose this deployment to simplify installation and configuration and/or avoid configuration changes to surrounding network devices.

A virtual wire is the default configuration, and should be used only when no switching or routing is needed. If you do not plan to use the default virtual wire, you should manually delete the configuration before proceeding with interface configuration to prevent it from interfering with other interface settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more interfaces. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when Layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

For more information on Layer 2 deployments, refer to the Layer 2 Networking Tech Note and/or the Securing Inter VLAN Traffic Tech Note.

Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

You must assign an IP address to each physical Layer 3 interface you configure. You can also create logical subinterfaces for each physical Layer 3 interface that allows you to segregate the traffic on the interface based on VLAN tag (when VLAN trunking is in use) or by IP address, for example for multi-tenancy.

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

The configuration example in this chapter illustrates how to integrate the firewall into your Layer 3 network using static routes. For information on other types of routing integrations, refer to the following documents:

- How to Configure OSPF Tech Note
- How to Configure BGP Tech Note

About Network Address Translation (NAT)

When you use private IP addresses within your internal networks, you must use network address translation (NAT) in order to translate the private addresses to public addresses that can be routed on external networks. In PAN-OS, you create NAT policy rules that instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For more details about the different types of NAT rules, refer to the Understanding and Configuring NAT Tech Note.

It is important to understand the way the firewall applies the NAT and security policies in order to determine what policies you need based on the zones you have defined. Upon ingress, the firewall inspects a packet to see if it matches any of the NAT rules that have been defined, based on source and/or destination zone. It then evaluates and applies any security rules that match the packet based on the original (pre-NAT) source and destination addresses. Finally, it translates the source and/or destination port numbers for any matching NAT rules upon egress. This distinction is important, because it means that the firewall determines what zone a packet is destined for based on the address on the packet, not in the placement of the device based on its internally assigned address.

About Security Policies

Security policies protect network assets from threats and disruptions and aid in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policies determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. By default, intra-zone traffic (that is traffic within the same zone, for example *from* trust *to* trust), is allowed. Traffic between different zones (or inter-zone traffic) is blocked until you create a security policy to allow the traffic.

Security policies are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

Components of a Security Policy

The security policy construct permits a combination of the required and optional components listed below.

	Field	Description
Required Fields	**Name**	A label that supports up to 31 characters, used to identify the rule.
	Source Zone	The zone from which the traffic originates.
	Destination Zone	The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone.
	Application	The application which you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed.
	Action	Specifies an *Allow* or *Deny* action for the traffic based on the criteria you define in the rule.
Optional Fields	**Tag**	A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a particular keyword or color, for example *Inbound to DMZ*.
	Description	A text field, up to 255 characters, used to describe the rule.
	Source Address	Define host IP or FQDN, subnet, named groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
	Destination Address	The location or destination for the traffic. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
	Source User	The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. See Configure User Identification for details.

	Field	Description (Continued)
	URL Category	Using the URL Category as match criteria allows you to customize security profiles (antivirus, anti-spyware, vulnerability, file-blocking, Data Filtering, and DoS) on a per-URL-category basis. For example, you can prevent the downloads/uploads of executable files for URL categories that represent higher risk while allowing them for other categories. This functionality also allows you to attach schedules to specific URL categories (allow social-media websites during lunch and after-hours), mark certain URL categories with QoS (financial, medical, and business), and select different log forwarding profiles on a per-URL-category-basis. Although you can manually configure URL categories on your device, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license. **Note** To just provide basic URL category filtering, define the URL Category as *Any* and attach a URL Filtering profile to the security policy. See Create Security Rules for information on using the default profiles in your security policy and see Control Access to Web Content for more details.
	Service	Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose *any*, specify a port, or use *application-default* to permit use of the standards-based port for the application. For example, for applications with well- known port numbers such as DNS, the *application-default* option will match against DNS traffic only on TCP port 53. You can also add a custom application and define the ports that the application can use. **Note** For inbound allow rules (for example, from untrust to trust), using *application-default* prevents applications from running on unusual ports and protocols. Application-default is the default option; while the device still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/protocols.
Optional Fields	**Security Profiles**	Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are only evaluated for rules that have an *allow* action. For more information, see Scan Traffic for Threats.
	HIP Profiles (for GlobalProtect)	Enables matching based on the state of the client system, such as whenever it has the latest patches and antivirus updates. For more information, refer to the section "Use Host Information in Policy Enforcement" in the *GlobalProtect Administrator's Guide*.
	Options	Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect.

Policy Best Practices

The task of safely enabling Internet access and preventing misuse of web access privileges, and exposure to vulnerabilities and attacks is a continuous process. The key principle when defining policy on the Palo Alto Networks firewall is to use a positive enforcement approach. Positive enforcement implies that you selectively allow what is required for day-to-day business operations as opposed to a negative enforcement approach where you would selectively block everything that is not allowed. Consider the following best practices when creating policy:

- If you have two or more zones with identical security requirements, combine them into one security rule.
- The ordering of rules is crucial to ensure the best match criteria. Because policy is evaluated top down, the more specific policy must precede the ones that are more general, so that the more specific rule is not *shadowed*. The term shadow refers to a rule that is not evaluated or is skipped because it is placed lower in the policy list. When the rule is placed lower, it is not evaluated because the match criteria was met by another rule that preceded it, thereby shadowing the rule from policy evaluation.
- To restrict and control access to inbound applications, in the security policy, explicitly define the port that the service/application will be listening on.
- Logging for broad allow rules—for example access to well known servers like DNS—can generate a lot of traffic. Hence it is not recommended unless absolutely necessary.
- By default, the firewall creates a log entry at the end of a session. However, you can modify this default behavior and configure the firewall to log at the start of the session. Because this significantly increases the log volume, logging at session start is recommended only when you are troubleshooting an issue. Another alternative for troubleshooting without enabling logging at session start is to use the session browser (**Monitor > Session Browser**) to view the sessions in real time.

About Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in a security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in a policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

You can create the following policy objects on the firewall:

Policy Object	Description
Address/Address Group, Region	Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an *address group* object. You can also use dynamic address groups to monitor changes and dynamically update IP addresses in environments where host IP addresses change frequently.
User/User Group	Allow you to create a list of users from the local database or an external database and group them. For information on creating user groups, refer to the User-ID section in the *PAN-OS Administrator's Guide*.
Application Group and Application Filter	An *Application Filter* allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can filter by one or more attributes—category, sub-category, technology, risk, characteristics—and save your application filter. With an application filter, when a PAN-OS content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter. An *Application Group* allows you to create a static group of specific applications that you wish to group together for a group of users or for a particular service.
Service/Service Groups	Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services—service-http and service-https—that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application). **Note** To view the standard ports used by an application, in **Objects > Applications** search for the application and click the link. A succinct description displays.

Some examples of *address* and *application* policy objects are shown in the security policies that are included in Create Security Rules. For information on the other policy objects, see Protect Your Network Against Threats and for more in-depth information refer to the *PAN-OS Administrator's Guide*.

About Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.

Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy

The different types of security profiles that can be attached to security policies are: Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking, and Data Filtering. The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Create Security Rules

for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See Scan Traffic for Threats for more information.

Set Up Interfaces and Zones

The following sections provide information on configuring interfaces and zones:

- Plan Your Deployment
- Configure Interfaces and Zones

Plan Your Deployment

Before you begin configuring your interfaces and zones, you should take some time to plan out the zones you will need based on the different usage requirements within your organization. In addition, you should gather all of the configuration information you will need ahead of time. At a basic level, you should plan which interfaces will belong to which zones. For Layer 3 deployments you'll also need to obtain the required IP addresses and network configuration information from your network administrator, including information on how to configure the routing protocol or static routes required for the virtual router configuration. The example in this chapter will be based on the following topology:

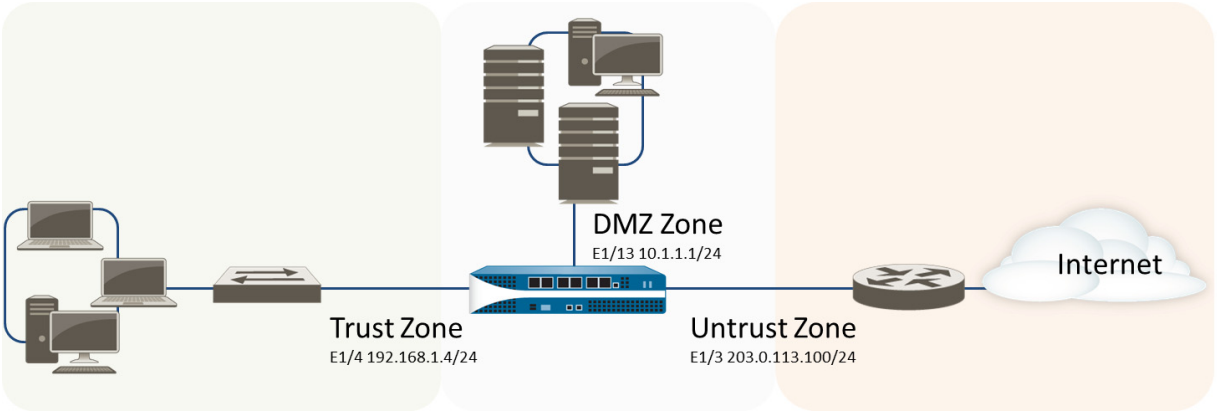

The following table shows the information we will use to configure the Layer 3 interfaces and their corresponding zones as shown in the sample topology.

Zone	Deployment Type	Interface(s)	Configuration Settings
Untrust	L3	Ethernet1/3	**IP address:** 203.0.113.100/24 **Virtual router:** default **Default route:** 0.0.0.0/0 **Next hop:** 203.0.113.1
Trust	L3	Ethernet1/4	**IP address:** 192.168.1.4/24 **Virtual router:** default
DMZ	L3	Ethernet1/13	**IP address:** 10.1.1.1/24 **Virtual router:** default

Configure Interfaces and Zones

After you plan your zones and the corresponding interfaces, you can configure them on the device. The way you configure each interface depends on your network topology.

The following procedure shows how to configure a Layer 3 deployment as depicted in the preceding example topology. For information on configuring a Layer 2 or virtual wire deployment, refer to the Networking section in the *PAN-OS Administrator's Guide*.

The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

Set Up Interfaces and Zones		
Step 1	Configure a default route to your Internet router.	1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog. 2. Click **Add** and select the **Static Routes** tab. Click **Add** on either the **IPv4** or **IPv6** tab, enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0). 3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address for your Internet gateway (for example, 203.0.113.1). 4. Click **OK** twice to save the virtual router configuration.
Step 2	Configure the external interface (the interface that connects to the Internet).	1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring ethernet1/3 as the external interface. 2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**. 3. In the **Virtual Router** drop-down, select **default**. 4. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**. 5. To assign an IP address to the interface, select the **IPv4** tab and **Static** radio button. Click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24. 6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**. 7. To save the interface configuration, click **OK**.

Set Up Interfaces and Zones (Continued)		
Step 3	Configure the interface that connects to your internal network. **Note** In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. See Configure NAT Policies.	1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface. 2. Select **Layer3** from the **Interface Type** drop down. 3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**. 4. Select the same Virtual Router you used in Step 2. 5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 7. To save the interface configuration, click **OK**.
Step 4	Configure the interface that connects to the DMZ.	1. Select the interface you want to configure. 2. Select **Layer3** from the **Interface Type** drop down. In this example, we are configuring Ethernet1/13 as the DMZ interface. 3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**. 4. Select the Virtual Router you used in Step 2, default in this example. 5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 7. To save the interface configuration, click **OK**.
Step 5	Save the interface configuration.	Click **Commit**.
Step 6	Cable the firewall.	Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment.
Step 7	Verify that the interfaces are active. 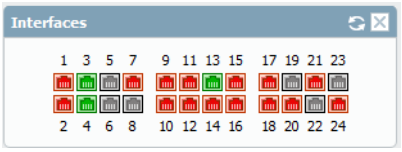	From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**.

Configure NAT Policies

Based on the example topology we used to create the interfaces and zones, there are three NAT policies we need to create as follows:

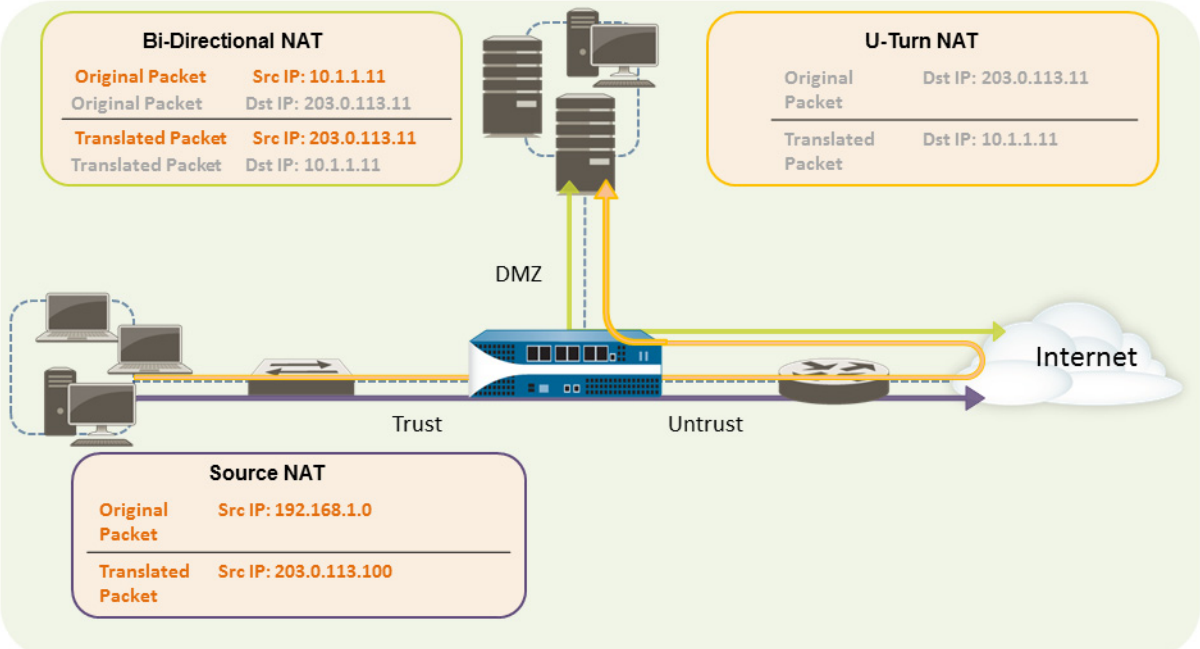

- To enable the clients on the internal network to access resources on the Internet, the internal 192.168.1.0 addresses will need to be translated to publicly routable addresses. In this case, we will configure source NAT, using the egress interface address, 203.0.113.100, as the source address in all packets that leave the firewall from the internal zone. See Translate Internal Client IP Addresses to your Public IP Address for instructions.
- To enable clients on the internal network to access the public web server in the DMZ zone, we will need to configure a NAT rule that redirects the packet from the external network, where the original routing table lookup will determine it should go based on the destination address of 203.0.113.11 within the packet, to the actual address of the web server on the DMZ network of 10.1.1.11. To do this you must create a NAT rule from the trust zone (where the source address in the packet is) to the untrust zone (where the original destination address is) to translate the destination address to an address in the DMZ zone. This type of destination NAT is called *U-Turn NAT*. See Enable Clients on the Internal Network to Access your Public Servers for instructions.
- To enable the web server—which has both a private IP address on the DMZ network and a public-facing address for access by external users—to both send and receive requests, the firewall must translate the incoming packets from the public IP address to the private IP address and the outgoing packets from the private IP address to the public IP address. On the firewall, you can accomplish this with a single bi-directional static source NAT policy. See Enable Bi-Directional Address Translation for your Public-Facing Servers.

Translate Internal Client IP Addresses to your Public IP Address

When a client on your internal network sends a request, the source address in the packet contains the IP address for the client on your internal network. If you use private IP address ranges internally, the packets from the client will not be able to be routed on the Internet unless you translate the source IP address in the packets leaving the network into a publicly routable address. On the firewall you can do this by configuring a source NAT policy that translates the source address and optionally the port into a public address. One way to do this is to translate the source address for all packets to the egress interface on your firewall as shown in the following procedure.

Configure Source NAT

Step 1 Create an address object for the external IP address you plan to use.

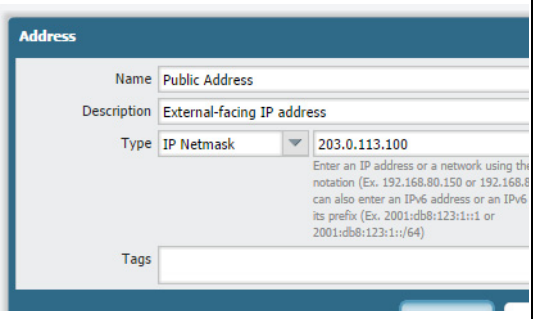

1. From the web interface, select **Objects > Addresses** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop down and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.
4. To save the address object, click **OK**.

Best practice:

Although you do not have to use address objects in your policies, it is a best practice because it simplifies administration by allowing you to make updates in one place rather than having to update every policy where the address is referenced.

Step 2 Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. Enter a **Name** for the policy and optional **Description**.
3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop down.

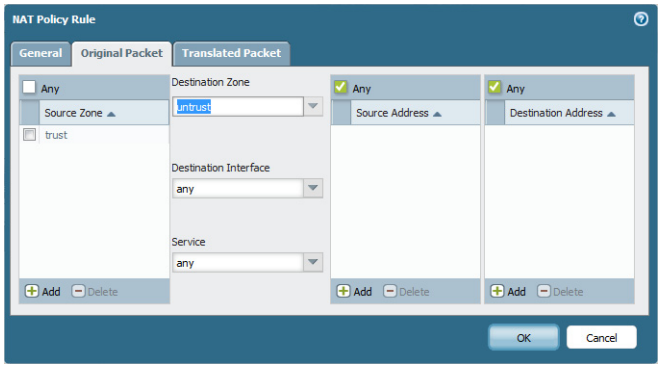

4. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** drop-down in the **Source Address Translation** section of the screen and then click **Add**. Select the address object you created in Step 1.

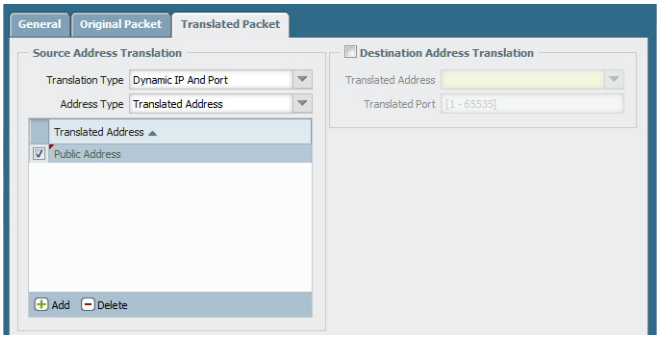

5. Click **OK** to save the NAT policy.

Configure Source NAT (Continued)	
Step 3 Save the configuration.	Click **Commit**.

Enable Clients on the Internal Network to Access your Public Servers

When a user on the internal network sends a request for access to the corporate web server in the DMZ, the DNS server will resolve to the public IP address. When processing the request, the firewall will use the original destination in the packet (the public IP address) and route the packet to the egress interface for the untrust zone. In order for the firewall to know that it must translate the public IP address of the web server to an address on the DMZ network when it receives requests from users on the trust zone, you must create a destination NAT rule that will enable the firewall to send the request to the egress interface for the DMZ zone as follows.

Configure U-Turn NAT

Step 1 Create an address object for the web server.

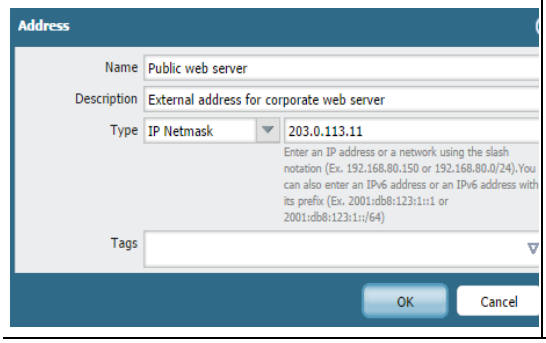
Address
Name
Public web server
Description
External address for corporate web server
Type
IP Netmask
203.0.113.11
Enter an IP address or a network using the slash notation (Ex. 192.168.80.150 or 192.168.80.0/24).You can also enter an IPv6 address or an IPv6 address with its prefix (Ex. 2001:db8:123:1::1 or 2001:db8:123:1::/64)
Tags
OK
Cancel

1. From the web interface, select **Objects > Addresses** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop-down and then enter the public IP address and netmask of the web server, 203.0.113.11/24 in this example.
4. To save the address object, click **OK**.

Step 2 Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. Enter a **Name** and optional **Description** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop down.

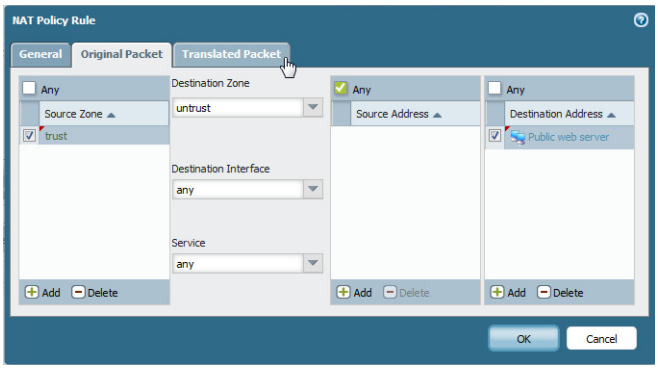

4. In the **Destination Address** section, click **Add** and select the address object you created for your public web server.

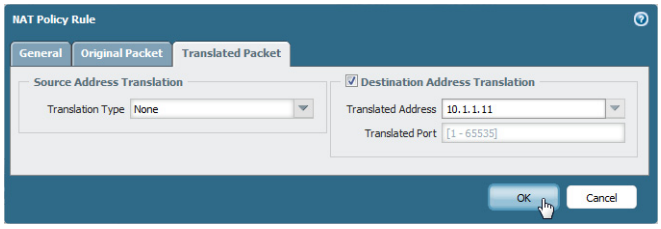

5. On the **Translated Packet** tab, select the **Destination Address Translation** check box and then enter the IP address that is assigned to the web server interface on the DMZ network, 10.1.1.11 in this example.
6. Click **OK** to save the NAT policy.

Configure U-Turn NAT (Continued)	
Step 3 Save the configuration.	Click **Commit**.

Enable Bi-Directional Address Translation for your Public-Facing Servers

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you will need a source NAT rule for translating the source address of the server to the external address upon egress. You do this by creating a static NAT rule that instructs the firewall to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example. However, in the case of a public-facing server, the server must both be able to send packets and receive them. In this case, you need a reciprocal policy that will translate the public address that will be the destination IP address in incoming packets from users on the Internet into the private address to enable the firewall to properly route the packet to your DMZ network. On the firewall you do this by creating a bi-directional static NAT policy as described in the following procedure.

Configure Bi-Directional NAT	
Step 1 Create an address object for the web server's internal IP address. 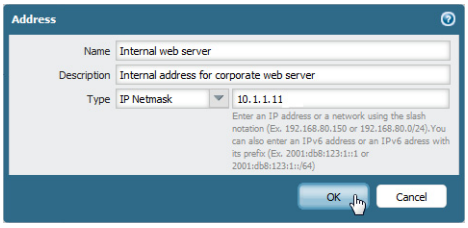	1. From the web interface, select **Objects > Addresses** and then click **Add**. 2. Enter a **Name** and optionally a **Description** for the object. 3. Select **IP Netmask** from the **Type** drop down and then enter the IP address and netmask of the web server on the DMZ network, 10.1.1.11/24 in this example. 4. To save the address object, click **OK**. **Note** If you did not already create an address object for the public address of your web server you should also create that object now.

Configure Bi-Directional NAT (Continued)

Step 2 Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. Enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your DMZ in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop down.
4. In the **Source Address** section, click **Add** and select the address object you created for your internal web server address.

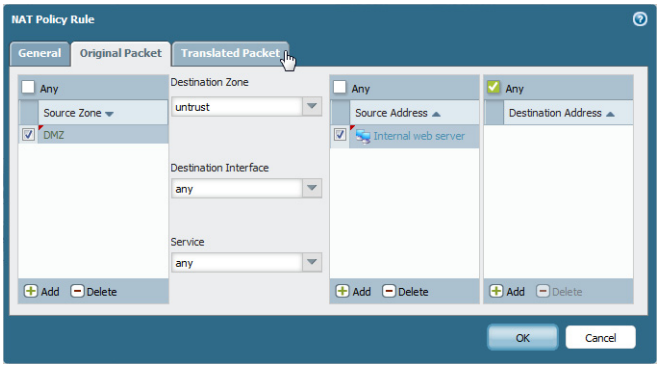

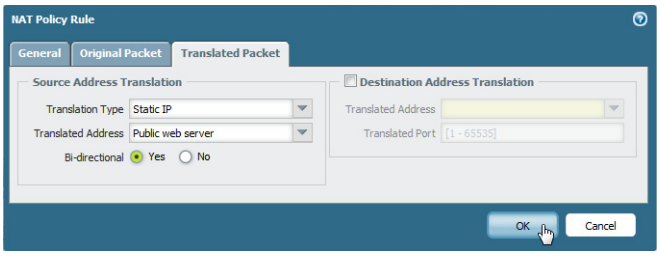

5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** drop down in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** drop down.
6. In the **Bi-directional** field, select **Yes**.
7. Click **OK** to save the NAT policy.

Step 3	Save the configuration.	Click **Commit**.

Set Up Basic Security Policies

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service, and Zone protection policies. Different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. The following topics describe basic security policies and the default security profiles:

- Create Security Rules
- Test Your Security Policies
- Monitor the Traffic on Your Network

This section describes security policies only. For information on the other types of policies see Protect Your Network Against Threats or refer to the following sections in *PAN-OS Administrator's Guide*: "Quality of Service" (for information on QoS policies) and "Map IP Addresses to User Names Using Captive Portal" (for information on Captive Portal policies).

Create Security Rules

Security policies reference security zones and enable you to allow, restrict, and track traffic on your network. Because each zone implies a level of trust, the implicit rule for passing traffic between two different zones is deny, and the traffic within a zone is permitted. To allow traffic between two different zones, you must create a security rule that allows traffic to flow between them.

While setting up the basic framework for securing the enterprise perimeter, it's a good idea to start with a simple security policy that allows traffic between the different zones without being too restrictive. As illustrated in the following section, our objective is to minimize the likelihood of breaking applications that users on the network need access to, while providing visibility into the applications and the potential threats for your network.

When defining policies make sure that you do not create a policy that denies all traffic from *any* source zone to *any* destination zone as this will break intra-zone traffic that is implicitly allowed. By default, intra-zone traffic is permitted because the source and destination zones are the same and therefore share the same level of trust.

Define Basic Security Rules

Step 1 Permit Internet access for all users on the enterprise network.

Zone: Trust to Untrust

Note By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone-naming convention.

To safely enable applications that are required for day-to-day business operations we will create a simple rule that allows access to the Internet. To provide basic threat protection, we will attach the default security profiles available on the firewall.

1. Select **Policies > Security** and click **Add.**
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, Set the **Destination Zone** to Untrust.

Note To scan policy rules and visually identify the zones on each rule, create a tag with the same name as the zone. For example, to color code the Trust zone as green, select **Objects > Tags,** click **Add** and **Name** the tag Trust, and select the **Color** green.

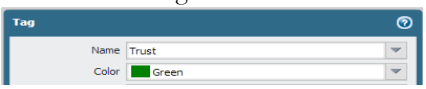

5. In the **Service/ URL Category** tab, select service-http and service-https.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting.**
7. Verify that logging is enabled at the end of a session under **Options**. Only traffic that matches a security rule will be logged.

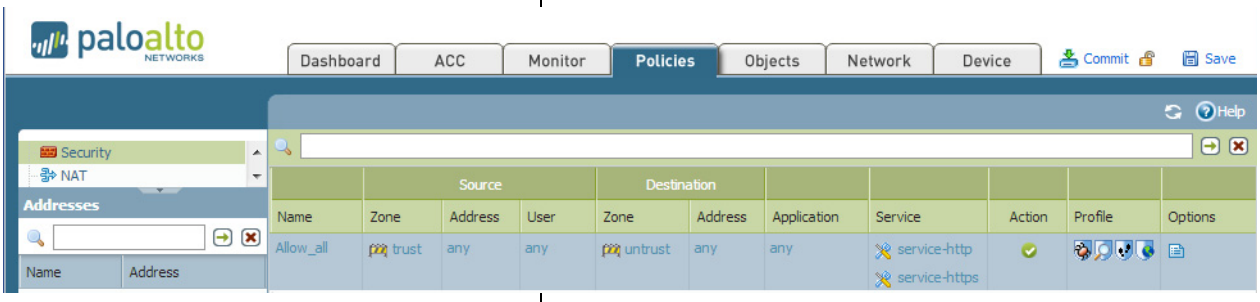

Step 2 Permit users on the internal network to access the servers in the DMZ.

Zone: Trust to DMZ

Note If using IP addresses for configuring access to the servers in the DMZ, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT addresses), and the post-NAT zone.

1. Click **Add** in the **Policies > Security** section.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, set the **Destination Zone** to DMZ.
5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.
6. In the **Actions** tab, set the **Action Setting** to Allow.
7. Leave all the other options at the default values.

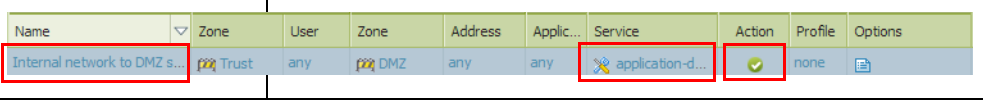

Define Basic Security Rules (Continued)

Step 3 Restrict access from the Internet to the servers on the DMZ to specific server IP addresses only.

For example, you might only allow users to access the webmail servers from outside.

Zone: Untrust to DMZ

To restrict inbound access to the DMZ from the Internet, configure a rule that allows access only to specific servers' IP addresses and on the default ports that the applications use.

1. Click **Add** to add a new rule, and give it a descriptive **Name**.
2. In the **Source** tab, set the **Source Zone** to Untrust.
3. In the **Destination** tab, set the **Destination Zone** to DMZ.
4. Set the **Destination Address** to the **Public web server** address object you created earlier. The public web server address object references the public IP address—208.80.56.11/24—of the web server that is accessible on the DMZ.
5. Select the webmail application in the **Application** tab.
6. Make sure the **Service** is set to **application-default**.

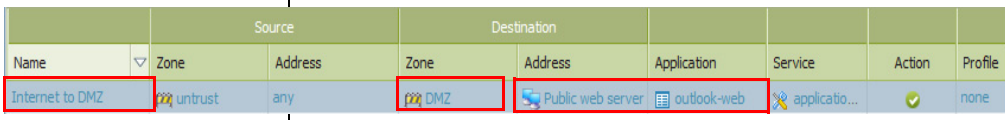

7. Set the **Action Setting** to **Allow**.

Step 4 Allow access from the DMZ to your internal network (Trust zone). To minimize risk, you will allow traffic only between specific servers and destination addresses. For example, if you have an application server on the DMZ that needs to communicate with a specific database server in your Trust zone, create a rule to allow traffic between a specific source to a specific destination.

Zone: DMZ to Trust

1. Click **Add** to add a new rule, and give it a descriptive name.
2. Set the **Source Zone** to DMZ.
3. Set the **Destination Zone** to Trust.
4. Create an address object that specifies the server(s) on your Trust zone that can be accessed from the DMZ.

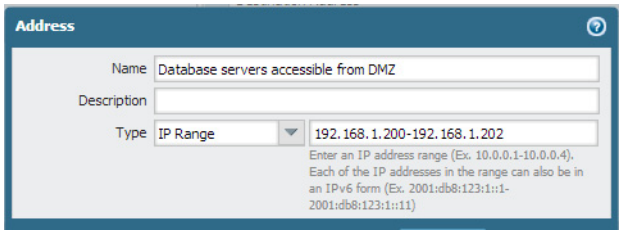
Address
Name
Database servers accessible from DMZ
Description
Type
IP Range
192.168.1.200-192.168.1.202
Enter an IP address range (Ex. 10.0.0.1-10.0.0.4). Each of the IP addresses in the range can also be in an IPv6 form (Ex. 2001:db8:123:1::1-2001:db8:123:1::11)

5. In the **Destination** tab on the Security Policy rule, set the **Destination Address** to the Address object you created above.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection, under **Profile Setting.**
   c. In the Other Settings section, select the option to **Disable Server Response Inspection**. This setting disables the antivirus and anti-spyware scanning on the server-side responses, and thus reduces the load on the firewall.

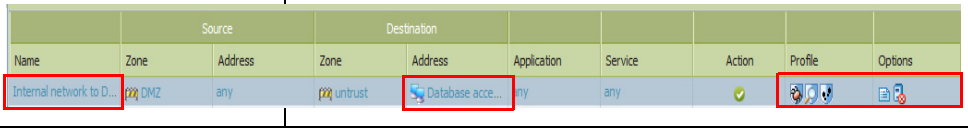

Define Basic Security Rules (Continued)		
Step 5	Enable the servers on the DMZ to obtain updates and hot fixes from the Internet. Say, for example, you would like to allow the Microsoft Update service. Zone: DMZ to Untrust	1. Add a new rule and give it a descriptive label. 2. Set the **Source Zone** to DMZ. 3. Set the **Destination Zone** to Untrust. 4. Create an application group to specify the applications that you would like to allow. In this example, we allow Microsoft updates (ms-updates) and dns. 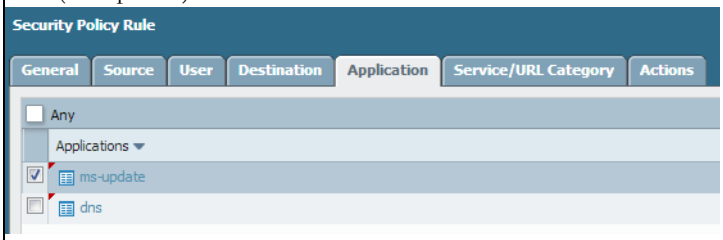 5. Ensure the **Service** is set to **application-default.** This allows the firewall to permit the applications only when they use the standard ports associated with these applications. 6. Set the **Action Setting** to **Allow**. 7. Attach the default profiles for antivirus, anti-spyware, and vulnerability protection, under **Profiles**. 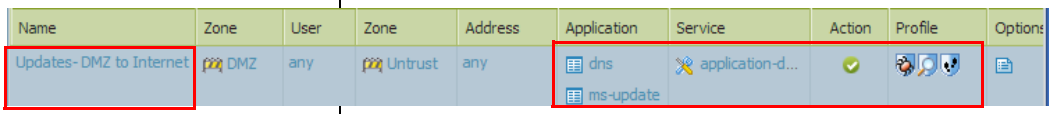
Step 6	Save your policies to the running configuration on the device.	Click **Commit**.

Test Your Security Policies

To verify that you have set up your basic policies effectively, test whether your security policies are being evaluated and determine which security rule applies to a traffic flow.

Verify Policy Match Against a Flow	

To verify the policy rule that matches a flow, use the following CLI command:
test security-policy-match source *<IP_address>* **destination** *<IP_address>* **destination port** *<port_number>* **protocol** *<protocol_number>*
The output displays the best rule that matches the source and destination IP address specified in the CLI command.
	For example, to verify the policy rule that will be applied for a server on the DMZ with the IP address 208.90.56.11 when it accesses the Microsoft update server, you will try the following command:
	`test security-policy-match source 208.80.56.11 destination 176.9.45.70 destination-port 80 protocol 6`
	`"Updates-DMZ to Internet" {`
	`from dmz;`
	`source any;`
	`source-region any;`
	`to untrust;`
	`destination any;`
	`destination-region any;`
	`user any;`
	`category any;`
	`application/service[ dns/tcp/any/53 dns/udp/any/53 dns/udp/any/5353 ms-update/tcp/any/80 ms-update/tcp/any/443];`
	`action allow;`
	`terminal yes;`

Monitor the Traffic on Your Network

Now that you have a basic security policy in place, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network, to identify where you need to create more granular policies.

Unlike traditional firewalls that use ports or protocols to identify applications, the Palo Alto Networks firewalls use the application signature (the App-ID technology) to monitor applications. The application signature is based on unique application properties and related transaction characteristics in combination with the port or protocol. Therefore, even when the traffic uses the right port/protocol, the firewall can deny access to content because the application signature is not a match. This feature allows you to safely enable applications by allowing parts of the application while blocking or controlling functions within the same application. For example, if you allow the application *web-browsing* a user will be able to access content on the Internet. Then, if a user goes to Facebook and then goes on to play Scrabble on Facebook, the firewall will identify the application shifts and recognize Facebook as an *application* and Scrabble as a *Facebook-app*. Therefore, if you create a specific rule that blocks Facebook applications, the user will be denied access to Scrabble while still being able to access Facebook.

To monitor traffic on your network:

- In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policies that block unwanted applications, while allowing and enabling applications in a secure manner.
- Determine what updates/modifications are required for your network security rules and implement the changes. For example:
  - Evaluate whether to allow content based on schedule, users, or groups
  - Allow or control certain applications or functions within an application
  - Decrypt and inspect content
  - Allow but scan for threats and exploits

    For information on refining your security policies and for attaching custom security profiles, see Protect Your Network Against Threats.
- View the traffic and threat logs at **Monitor > Logs**.

Traffic logs are dependent on how your security policies are defined and setup to log traffic. The **ACC** tab, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter-zone traffic that is allowed by policy and the intra-zone traffic that is allowed implicitly.

- Review the URL filtering logs to scan through alerts, denied categories/URLs. In order to generate a URL log, you must have a URL profile attached to the security rule and the action must be set to alert, continue, override or block.

Protect Your Network Against Threats

The Palo Alto Networks next-generation firewall has unique threat prevention capabilities that allow it to protect your network from attack despite evasive, tunneled, or circumvention techniques. The threat prevention features on the firewall include the WildFire service, the Security Profiles that support Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking and Data Filtering capabilities and the Denial of Service (DoS) and Zone protection functionality.

Before you can apply threat prevention features, you must first configure zones—to identify one or more source or destination interfaces—and security policies. To configure interfaces, zones, and the policies that are needed to apply threat prevention features, see Set Up Interfaces and Zones and Set Up Basic Security Policies.

To begin protecting your network from threats start here:

- Enable WildFire
- Scan Traffic for Threats
- Control Access to Web Content

Enable WildFire

The WildFire service is included as part of the base product. The WildFire service enables the firewall to forward attachments to a sandbox environment where applications are run to detect any malicious activity. As new malware is detected by the WildFire system, malware signatures are automatically generated and are made available within 24-48 hours in the antivirus daily downloads. Your threat prevention subscription entitles you for antivirus signature updates that include signatures discovered by WildFire.

Consider purchasing the WildFire subscription service for these additional benefits:

- Sub-hourly (as often as every 15 minutes) WildFire signature updates
- Advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet)
- Ability to upload files using the WildFire API
- Ability to forward files to a private WF-500 WildFire appliance

  While the ability to configure a file blocking profile to forward Portable Executable (PE) files to the WildFire cloud for analysis is free, in order to forward files to a private WildFire appliance, a WildFire subscription is required.

Enable WildFire		
Step 1	Confirm that your device is registered and that you have a valid support account as well as any subscriptions you require.	1. Go to the Palo Alto Networks Support Site, log in, and select **My Devices**. 2. Verify that the firewall is listed. If it is not listed, see Register With Palo Alto Networks. 3. (Optional) Activate Licenses.
Step 2	Set the WildFire forwarding options.	1. Select **Device > Setup > WildFire**. 2. Click the edit icon in the General Settings section. 3. (Optional) Specify the **WildFire Server** to which to forward files. By default, the firewall will forward files to the public WildFire cloud hosted in the United States. To forward files to a different WildFire cloud, enter a new value as follows: • To forward to a private WildFire cloud, enter the IP address or FQDN of your WF-500 WildFire appliance. • To forward files to the public WildFire cloud running in Japan, enter `wildfire.paloaltonetworks.jp`. 4. (Optional) If you want to change the maximum file size that the firewall can forward for a specific type of file, modify the value in the corresponding field. 5. Click **OK** to save your changes.
Note	If you do not have a WildFire subscription you can only forward executables.	

Enable WildFire (Continued)		
Step 3	Set up a file blocking profile to forward files to WildFire.	1. Select **Objects > Security Profiles > File Blocking** and click **Add**. 2. Enter a **Name** and optionally a **Description** for the profile. 3. Click **Add** to create a forwarding rule and enter a name. 4. In the **Action** column, select **forward**. 5. Leave the other fields set to **any** to forward any supported file type from any application. 6. Click **OK** to save the profile.
Step 4	Attach the file blocking profile to the security policies that allow access to the Internet.	1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules. 2. Click the **Actions** tab within the security policy. 3. In the Profile Settings section, click the drop-down and select the file blocking profile you created for WildFire forwarding. (If you don't see a drop-down for selecting a profile, select **Profiles** from the **Profile Type** drop-down.
Step 5	Save the configuration.	Click **Commit**.
Step 6	Verify that the firewall is forwarding files to WildFire.	1. Select **Monitor > Logs > Data Filtering**. 2. Check the **Action** column for the following actions: • **Forward**— Indicates that the file was successfully forwarded by the file blocking profile attached to the security policy. • **Wildfire-upload-success**—Indicates that the file was sent to WildFire. This means the file is not signed by a trusted file signer and it has not been previously analyzed by WildFire. • **Wildfire-upload-skip**—Indicates that the file was identified as eligible to be sent to WildFire by a file blocking profile/security policy, but did not need to be analyzed by WildFire because it has already been analyzed previously. In this case, the action will display as **forward** in the Data Filtering log because it was a valid forward action, but it was not sent to WildFire and analyzed because the file has already been sent to the WildFire cloud from another session, possibly from another firewall. 3. View the WildFire logs by selecting **Monitor > Logs > WildFire Submissions**. If new WildFire logs appear, the firewall is successfully forwarding files to WildFire and WildFire is returning file analysis reports.

For more information on WildFire, refer to the Palo Alto Networks WildFire Administrator's Guide.

Scan Traffic for Threats

Security profiles provide threat protection in security policies. For example, you can apply an antivirus profile to a security policy and all traffic that matches the security policy will be scanned for viruses.

The following sections provide steps for setting up a basic threat prevention configuration:

- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
- Set Up File Blocking

Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with default Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to security policies.

As a best practice, attach the default antivirus profile and the strict profiles for anti-spyware and vulnerability protection, to your basic web access policies to ensure that the traffic entering your network is free from threats.

As you monitor the traffic on your network and expand your policy rulebase, you can design more granular profiles to address your specific security needs. The strict profile applies the block response to all client and server critical, high and medium severity spyware events and uses the default action for low and informational vulnerability protection events. For all Anti-Spyware and Vulnerability Protection signatures a default action is defined by Palo Alto Networks. You can view the default action by navigating to **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles >Vulnerability Protection** and then selecting a profile. Click the **Exceptions** tab and then click **Show all signatures** and you will see a list of the signatures with the default action in the **Action** column. To change the default action, you must create a new profile and then create rules with a non-default action, and/or add individual signature exceptions in the **Exceptions** tab of the profile.

Set up Antivirus/Anti-Spyware/Vulnerability Protection		
Step 1	Verify that you have a Threat Prevention license.	• The Threat Prevention license bundles the Antivirus, Anti-Spyware, and the Vulnerability Protection features in one license. • Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and valid (check the expiration date).
Step 2	Download the latest antivirus threat signatures.	1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures. 2. In the **Actions** column, click **Download** to install the latest Antivirus, and Applications and Threats signatures.

Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)

Step 3 Schedule signature updates.

Best Practice for Updates:

Perform a **download-and-install** on a daily basis for antivirus updates and weekly for applications and threats updates.

1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.
2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.
3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.
4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device.

Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)	
Step 4 Attach the security profiles to a security policy.	1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab. 2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus,** and strict for **Vulnerability Protection,** and **Anti-Spyware**. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.) 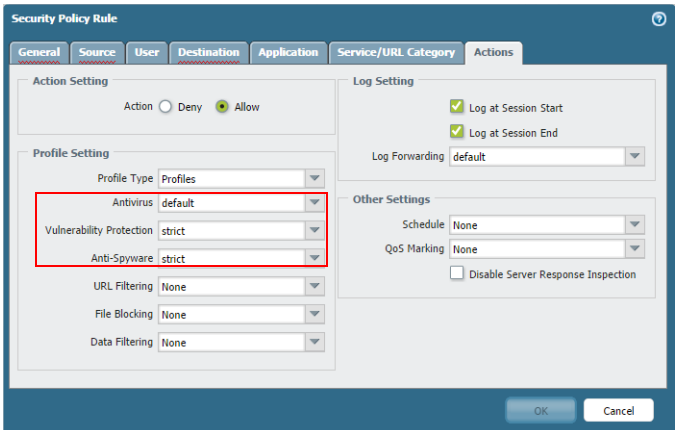
Step 5 Save the configuration.	Click **Commit**.

Set Up File Blocking

File blocking profiles allow you to identify specific file types that you want to want to block or monitor. The following workflow shows how to set up a basic file blocking profile that prevents users from downloading executable files from the Internet.

Configure File Blocking	
Step 1 Create the file blocking profile.	1. Select **Objects > Security Profiles > File Blocking** and click **Add**. 2. Enter a **Name** for the file blocking profile, for example *Block_EXE*. Optionally enter a **Description**, such as *Block users from downloading exe files from websites.*

Configure File Blocking (Continued)	
Step 2 Configure the file blocking options.	1. Click **Add** to define the profile settings. 2. Enter a **Name**, such as *BlockEXE*. 3. Set the **Applications** to which to apply file blocking, or leave it set to any. 4. Set **File Types** to block. For example, to block download of executables, you would select **exe**. 5. Specify the **Direction** in which to block files **download**, **upload**, or **both**. 6. Set the **Action** to one of the following: • **continue**—Users will have to click Continue in order to proceed with the download/upload. You must enable response pages on the associated interfaces if you plan to use this option. • **block**—Files matching the selected criteria will be blocked from download/upload. • **alert**—Files matching the selected criteria will be allowed, but will generate a log entry in the data filtering log. 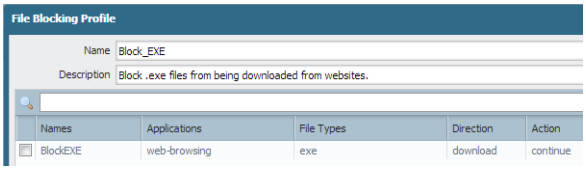 7. Click **OK** to save the profile.
Step 3 Attach the file blocking profile to the security policies that allow access to content.	1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules. 2. Click the **Actions** tab within the security policy. 3. In the Profile Settings section, click the drop-down and select the file blocking profile you created. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.)
Step 4 Enable Response Pages in the management profile for each interface on which you are attaching file blocking profile with a **continue** action.	1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile. 2. Select **Response Pages**, as well as any other management services required on the interface. 3. Click **OK** to save the interface management profile. 4. Select **Network > Interfaces** and select the interface to which to attach the profile. 5. On the **Advanced > Other Info** tab, select the interface management profile you just created. 6. Click **OK** to save the interface settings.

Configure File Blocking (Continued)

Step 5 To test the file blocking configuration, access a client PC in the trust zone of the firewall and attempt to download a .exe file from a website in the untrust zone. A response page should display. Click **Continue** to download the file. You can also set other actions, such as alert only, forward (which will forward to WildFire), or block, which will not provide a continue page to the user. The following shows the default response page for File Blocking:

Example: Default File Blocking Response Page

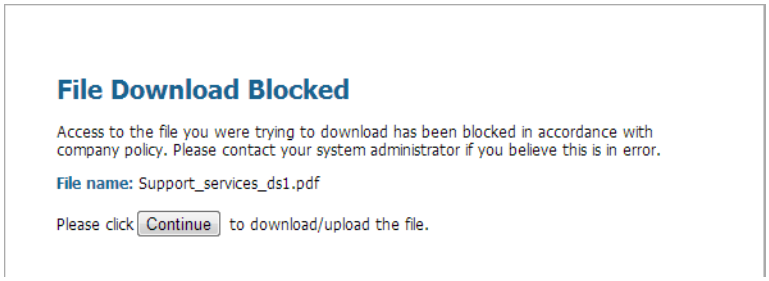
File Download Blocked
Access to the file you were trying to download has been blocked in accordance with company policy. Please contact your system administrator if you believe this is in error.
File name: Support_services_ds1.pdf
Please click Continue to download/upload the file.

Control Access to Web Content

URL filtering provides visibility and control over web traffic on your network. With URL filtering enabled, the firewall can categorize web traffic into one or more (from approximately 60) categories. You can then create policies that specify whether to allow, block, or log (alert) traffic based on the category to which it belongs. The following workflow shows how to enable PAN-DB for URL filtering, create security profiles, and attach them to security policies to enforce a basic URL filtering policy.

Configure URL Filtering	
Step 1 Confirm license information for URL Filtering.	1. Obtain and install a URL Filtering license. See Activate Licenses for details. 2. Select **Device > Licenses** and verify that the URL Filtering license is valid.
Step 2 Download the seed database and activate the license.	1. To download the seed database, click **Download** next to **Download Status** in the PAN-DB URL Filtering section of the Licenses page. 2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download. 3. After the download completes, click **Activate**. 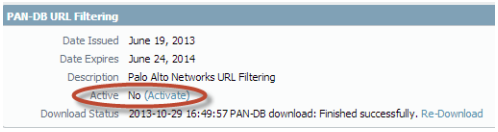
Step 3 Create a URL filtering profile. Best Practice for New Profiles: Because the default URL filtering profile blocks risky and threat-prone content, clone this profile when creating a new profile in order to preserve the default settings.	1. Select **Objects > Security Profiles > URL Filtering**. 2. Select the default profile and then click **Clone**. The new profile will be named *default-1*. 3. Select the new profile and rename it.

Configure URL Filtering (Continued)

Step 4 Define how to control access to web content. If you are not sure what traffic you want to control, consider setting the categories (except for those blocked by default) to alert. You can then use the visibility tools on the firewall, such as the ACC and App Scope, to determine which web categories to restrict to specific groups or to block entirely. You can then go back and modify the profile to block and allow categories as desired. You can also define specific sites to always allow or always block regardless of category and enable the safe search option to filter search results when defining the URL filtering profile. For more details, refer to "URL Filtering" in the PAN-OS Administrator's Guide.	1. For each category that you want visibility into or control over, select a value from the **Action** column as follows: • If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **Allow**. • For visibility into traffic to sites in a category, select **Alert** • To prevent access to traffic that matches the associated policy, select **Block** (this also generates a log entry). 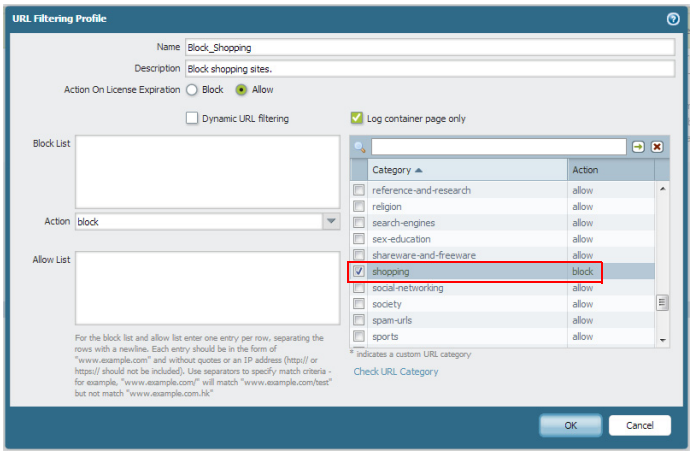 2. Click **OK** to save the URL filtering profile.
Step 5 Attach the URL filtering profile to a security policy.	1. Select **Policies > Security**. 2. Select the desired policy to modify it and then click the **Actions** tab. 3. If this is the first time you are defining a security profile, select **Profiles** from the **Profile Type** drop-down. 4. In the **Profile Settings** list, select the profile you just created from the **URL Filtering** drop-down. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.) 5. Click **OK** to save the profile. 6. **Commit** the configuration.
Step 6 Enable Response Pages in the management profile for each interface on which you are filtering web traffic.	1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile. 2. Select **Response Pages**, as well as any other management services required on the interface. 3. Click **OK** to save the interface management profile. 4. Select **Network > Interfaces** and select the interface to which to attach the profile. 5. On the **Advanced > Other Info** tab, select the interface management profile you just created. 6. Click **OK** to save the interface settings.

Configure URL Filtering (Continued)		
Step 7	Save the configuration.	Click **Commit**.
Step 8	To test URL filtering, access a client PC from the zone where the security policy is applied and attempt to access a site in a blocked category. You should see a URL Filtering response page that indicates that the page has been blocked:	

Web Page Blocked
Access to the web page you were trying to visit has been blocked in accordance with company policy. Please contact your system administrator if you believe this is in error.
User: 192.168.2.10
URL: amazon.com/
Category: shopping

For More Information

For more details on URL filtering, refer to the URL Filtering section of the PAN-OS Administrator's Guide.

For filtering or scanning for threats within encrypted traffic, such as with SSL and SSH, you must configure a decryption policy, see Decryption in the PAN-OS Administrator's Guide for more details.

For information about the threats and applications that Palo Alto Networks products can identify, visit the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

Configure User Identification

User Identification (User-ID) is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses. The following sections describe the Palo Alto Networks User-ID feature and provide instructions on setting up user- and group-based access:

- User Identification Overview
- Enable User Identification
- Enable User- and Group-Based Policy
- Verify the User-ID Configuration

User Identification Overview

User-ID seamlessly integrates Palo Alto Networks firewalls with a range of enterprise directory and terminal services offerings, enabling you to tie application activity and security policies to users and groups—not just IP addresses. The Palo Alto Networks next-generation firewall supports the following enterprise services:

- Microsoft Active Directory
- LDAP
- Novell eDirectory
- Citrix Metaframe Presentation Server or XenApp
- Microsoft Terminal Services

To be able to create policy based on user and group, the firewall must have a list of all available users and their corresponding group mappings that you can select from when defining your policies. It gets this group mapping information by connecting directly to your LDAP directory server. See About Group Mapping for more information.

To be able to enforce the user- and group-based policies, the firewall must be able to map the IP addresses in the packets it receives into user names. It gets this user mapping information—either directly or from a User-ID agent installed on a Windows server—by monitoring Microsoft Exchange Server or domain controller event logs for logon events, monitoring Novell eDirectory for login information, or by directly probing the client systems. See About User Mapping for more information.

About Group Mapping

In order to define security policies based on user or group, the firewall must retrieve the list of groups and the corresponding list of members from your directory server. To enable this functionality, you must create an LDAP server profile that instructs the firewall how to connect and authenticate to the server and how to search the directory for the user and group information. After you successfully connect to the LDAP server and configure the group mapping functionality for user identification, you will be able to select users or groups when defining your security policies. The firewall supports a variety of LDAP directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server.

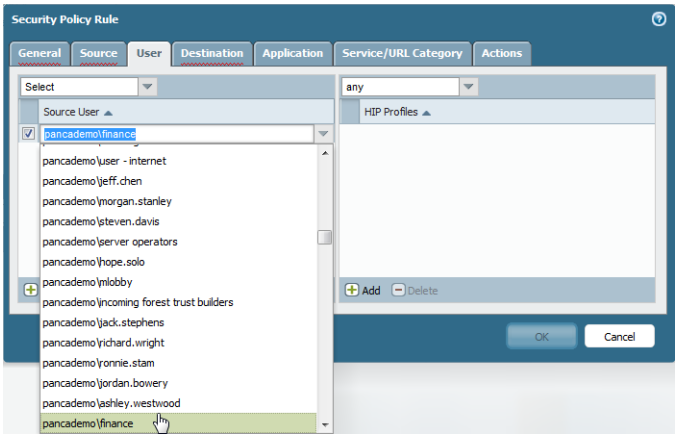

About User Mapping

Having the names of the users and groups is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security policies can be enforced appropriately. To do this user mapping, you must configure a User-ID agent—either by installing the agent software on a Windows server in the domain or by enabling the native agent on the firewall—to obtain the mappings using one or more of the following methods:

- Monitoring the security event logs for your Microsoft Exchange Servers, domain controllers, or Novell eDirectory servers for logon events. For example, in an AD environment, the agent will monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections (for monitored servers). Keep in mind that in order for these events to be recorded in the security log, the AD domain must be configured to log successful account logon events. See Configure User Mapping for details.
- In a Microsoft Windows environment, the agent can be configured to probe client systems using Windows Management Instrumentation (WMI) or NetBIOS (not supported on the PAN-OS integrated User-ID agent). If probing is enabled, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping it will send the address to the agent for an immediate probe. See Configure User Mapping for details.
- In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Services Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Services Agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. Refer to "Configure User Mapping for Terminal Server Users" in the PAN-OS Administrator's Guide for configuration details.
- In environments with existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—the firewall User-ID agent (either the Windows agent or the PAN-OS integrated agent on the firewall) can now listen for authentication syslog messages from those services. Syslog filters, which are provided by a content update (integrated User-ID agent only) or configured manually, allow the User-ID agent to parse and extract usernames and IP addresses from authentication syslog events generated by the external service, and add the information to the User-ID IP address to username mappings maintained by the firewall. Refer to "Configure User-ID to Receive User Mappings from a Syslog Sender" in the PAN-OS Administrator's Guide for configuration details.
- For mobile or roaming users, the GlobalProtect client provides the user mapping information to the firewall directly. For more information on setting up GlobalProtect, refer to For more information on setting up GlobalProtect, refer to the GlobalProtect Administrator's Guide.
- If the firewall or the User-ID agent are unable to map an IP address to a user—for example if the user is not logged in or is using an operating system such as Linux that is not supported by your domain servers—you can configure Captive Portal. When configured, any web traffic (HTTP or HTTPS) matching your Captive

Portal policy requires user authentication, either transparently via an NT LAN Manager (NTLM) challenge to the browser, or actively by redirecting the user to a web authentication form for authentication against a RADIUS, LDAP, Kerberos, or local authentication database or using client certificate authentication.

- For other types of user access that cannot be mapped using standard user mapping or Captive Portal methods—for example, to add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x enabled wireless network. Refer to the PAN-OS XMLAPI Usage Guide.

The following diagram illustrates the different methods that are used to identify users and groups on your network:

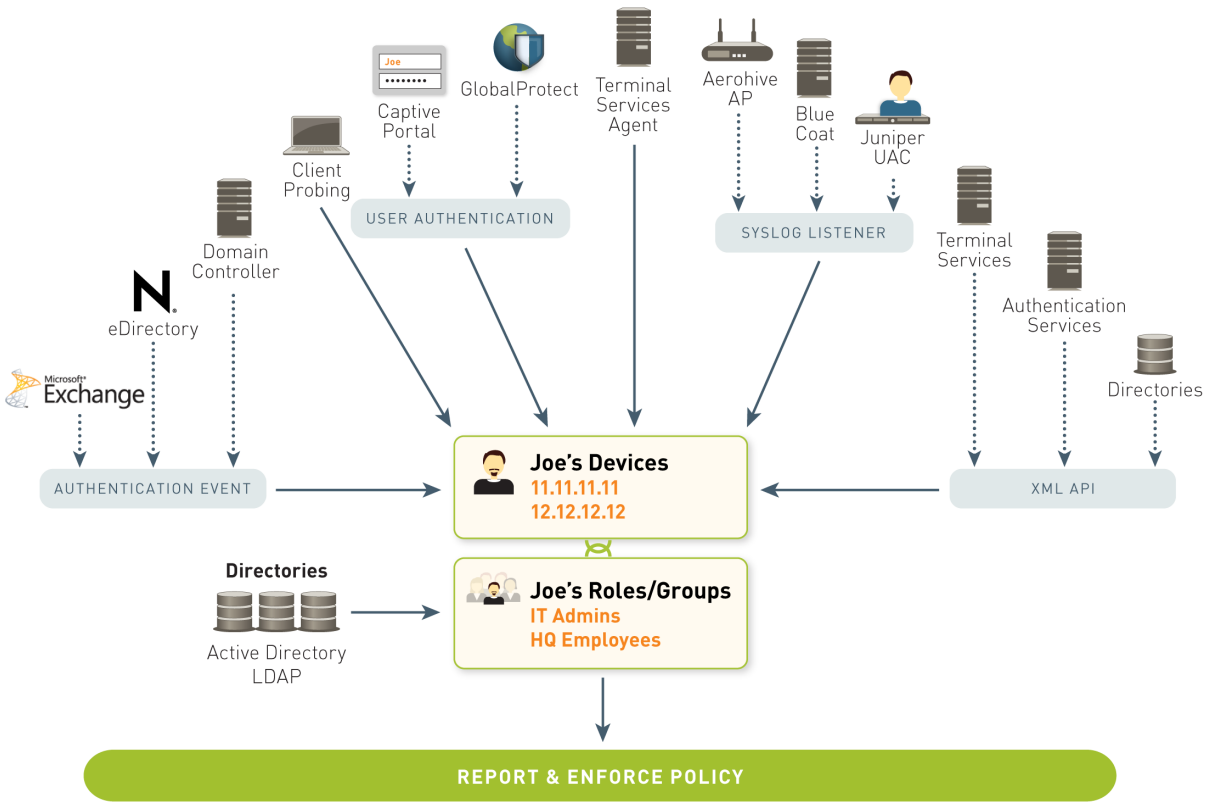

Enable User Identification

To enable policy enforcement based on user or group, you must complete the following tasks:

- Map Users to Groups
- Map IP Addresses to Users
- Enable User- and Group-Based Policy

Map Users to Groups

Use the following procedure to connect to your LDAP directory to enable the firewall to retrieve user-to-group mapping information:

Map Users to Groups

Step 1 Create an LDAP Server Profile that specifies how to connect to the directory servers you want the firewall to use to obtain group mapping information.

1. Select **Device > Server Profiles > LDAP**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
4. Click **Add** to add a new LDAP server entry and then enter a **Server** name to identify the server (1-31 characters) and the IP **Address** and **Port** number the firewall should use to connect to the LDAP server (default=389 for LDAP; 636 for LDAP over SSL). You can add up to four LDAP servers to the profile, however, all the servers you add to a profile must be of the same type. For redundancy you should add at least two servers.

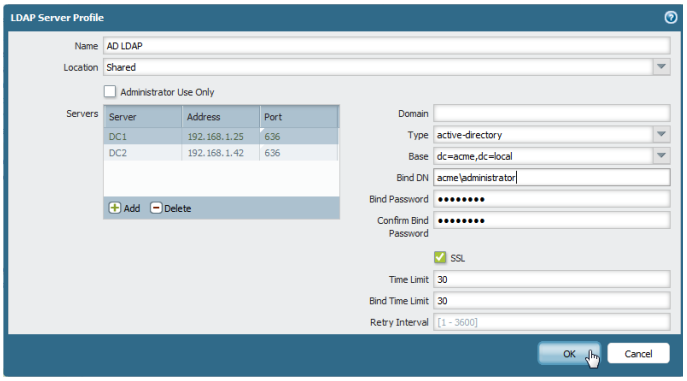

5. Enter the LDAP **Domain** name to prepend to all objects learned from the server. The value you enter here depends on your deployment:
   - If you are using Active Directory, you must enter the NetBIOS domain name; NOT a FQDN (for example, enter `acme`, not `acme.com`). Note that if you need to collect data from multiple domains you will need to create separate server profiles.
   - If you are using a global catalog server, leave this field blank.
6. Select the **Type** of LDAP server you are connecting to. The correct LDAP attributes in the group mapping settings will automatically be populated based on your selection. However, if you have customized your LDAP schema you may need to modify the default settings.
7. In the **Base** field, select the DN that corresponds to the point in the LDAP tree where you want the firewall to begin its search for user and group information.
8. Enter the authentication credentials for binding to the LDAP tree in the **Bind DN**, **Bind Password**, and **Confirm Bind Password** fields. The Bind DN can be in either User Principal Name (UPN) format (i.e. `administrator@acme.local`) or it can be a fully qualified LDAP name (i.e. `cn=administrator,cn=users,dc=acme,dc=local`).
9. If you want the firewall to communicate with the LDAP server(s) over a secure connection, select the SSL check box. If you enable SSL, make sure that you have also specified the appropriate port number.

Map Users to Groups (Continued)

Step 2 Add the LDAP server profile to the User-ID Group Mapping configuration.

1. Select **Device > User Identification > Group Mapping Settings** and click **Add**.
2. Select the **Server Profile** you created in Step 1.
3. Make sure the **Enabled** check box is selected.
4. (Optional) If you want to limit which groups are displayed within security policy, select the **Group Include List** tab and then browse through the LDAP tree to locate the groups you want to be able to use in policy. For each group you want to include, select it in the **Available Groups** list and click the add icon to move it to the **Included Groups** list. Repeat this step for every group you want to be able to use in your policies.
5. Click **OK** to save the settings.

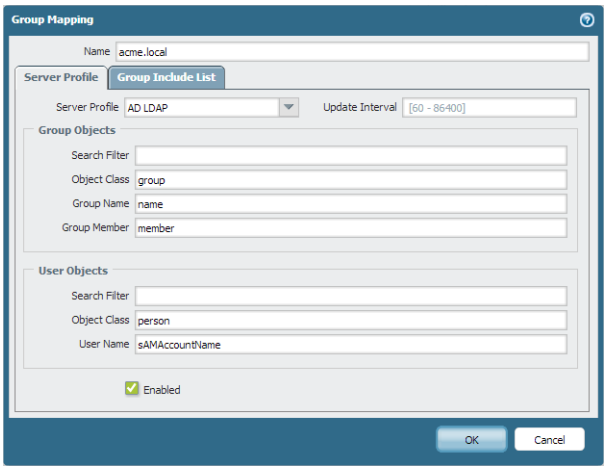

Step 3	Save the configuration.	Click **Commit**.

Map IP Addresses to Users

The tasks you need to perform to map IP addresses to user names depends on the type and location of the client systems on your network. Complete as many of the following tasks as necessary to enable mapping of your client systems:

- To map clients that are logged in to your monitored Exchange servers, domain controllers, or eDirectory servers, or Windows clients that can be directly probed see Configure User Mapping.
- If you have users with client systems that are not logged into your domain servers—for example, users running Linux clients that do not log in to the domain—see Map IP Addresses to User Names Using Captive Portal.
- If you have clients running multi-user systems such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, refer "Configure User Mapping for Terminal Server Users" in the PAN-OS Administrator's Guide for configuration details.
- For other clients that you are unable to map using the previous methods, you can use the XML-based REST API to add user mappings directly to the firewall. Refer to the PAN-OS XML API Usage Guide for instructions.

Configure User Mapping

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID Agent will perform the IP address to user mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your directory servers.

As a best practice, locate your User-ID agents on or near your monitored servers. This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of IP address mappings since the last update—from the agent to the firewall.

However, in smaller environments (although this will vary depending on deployment, as a general rule you should use the on-device agents in environments where you will be monitoring ten or fewer directory servers), you can use the on-device agent that resides on the firewall without the need to install separate agent software on your network servers. In addition, if you are using the on-device agent, you can configure it to redistribute user mapping information to other firewalls.

For information about the system requirements for installing the Windows-based User-ID agent, refer to the *User-ID Agent Release Notes*, which are available on the Palo Alto Networks Software Updates page.

The following procedure shows how to configure the PAN-OS integrated User-ID agent on the firewall for monitoring Active Directory domain controllers. For instructions on installing and configuring the Windows-based User-ID agent, refer to "Configure User Mapping Using the Windows User-ID Agent" in the PAN-OS Administrator's Guide.

Map IP Addresses to Users Using the PAN-OS Integrated User-ID Agent		
Step 1	Create an Active Directory account for the firewall agent that has the privilege levels required to log in to each service or host you plan to monitor to collect user mapping data.	• **Windows 2008 or later domain servers**—Add the account to the Event Log Readers group. If you are using the on-device User-ID agent, the account must also be a member of the Distributed COM Users Group. • **Windows 2003 domain servers**—Assign Manage Auditing and Security Logs permissions through group policy. • **WMI probing**—Make sure the account has rights to read the CIMV2 namespace; by default, Domain Administrator and Server Operator accounts have this permission. • **NTLM authentication**—Because the firewall must join the domain if you are using NTLM authentication with an on-device User-ID agent, the Windows account you create for NTLM access must have administrative privileges. Note that due to AD restrictions on virtual systems running on the same host, if you have configured multiple virtual systems, only vsys1 will be able to join the domain.

Map IP Addresses to Users Using the PAN-OS Integrated User-ID Agent (Continued)

Step 2	Define the servers the firewall should monitor to collect IP address to user mapping information. You can define entries for up to 100 Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory servers on your network. Keep in mind that in order to collect all of the required mappings, you must connect to all servers that your users log in to so that the firewall can monitor the security log files on all servers that contain logon events.	1. Select **Device > User Identification > User Mapping**. 2. In the **Server Monitor** section of the screen, click **Add**. 3. Enter a **Name** and **Network Address** for the server. The network address can be a FQDN or an IP address. 4. Select the **Type** of server. 5. Make sure the **Enabled** check box is selected and then click **OK** 6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Discover**. 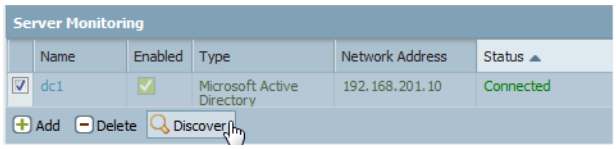 **Note** The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor. 7. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, in the **Palo Alto Networks User ID Agent Setup** section of the screen, click the Edit icon and then select the **Server Monitor** tab. Modify the value in the **Server Log Monitor Frequency (sec)** field. Best Practice: You should increase the value in the **Server Log Monitor Frequency (sec)** field to 5 seconds in environments with older DCs or high-latency links. 8. Click **OK** to save the changes.
Step 3	Set the domain credentials for the account the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing.	1. Click the Edit icon in the **Palo Alto Networks User ID Agent Setup** section of the screen. 2. On the **WMI Authentication** tab, enter the **User Name** and **Password** for the account that will be used to probe the clients and monitor servers. Enter the user name using the domain\username syntax. 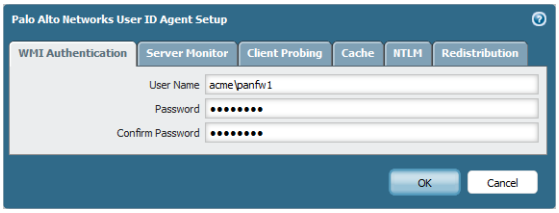

Map IP Addresses to Users Using the PAN-OS Integrated User-ID Agent (Continued)		
Step 4	(Optional) Enable WMI probing. **Note** The on-device agent does not support NetBIOS probing; it is supported on the Windows-based User-ID agent only.	1. On the **Client Probing** tab, select the **Enable Probing** check box. 2. (Optional) If necessary, modify the value of the **Probe Interval** to ensure that it is long enough for all learned IP addresses to be probed. 3. Make sure the Windows firewall will allow client probing by adding a a remote administration exception to the Windows firewall for each probed client.
Step 5	Save the configuration.	1. Click **OK** to save the User-ID agent setup settings. 2. Click **Commit** to save the configuration.

Step 6	(Optional) Define the set of users for which you do not need to provide IP address to user name mappings, such as service accounts or kiosk accounts.
	Note You can also use the `ignore-user` list to identify users that you want to force to authenticate using Captive Portal.
		1. Open a CLI session to the firewall.
		2. To add the list of user accounts for which you do not want the firewall to perform mapping, run the following command:
		`set user-id-collector ignore-user <value>`
		Where `<value>` is a list of the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Separate entries with a space and do not include the domain name with the username. For example:
		`set user-id-collector ignore-user SPAdmin SPInstall TFSReport`
		3. **Commit** your changes.

Step 7	Verify the configuration.	1. From the CLI, enter the following command: `show user server-monitor state all` 2. On the **Device > User Identification > User Mapping** tab in the web interface, verify that the **Status** of each server you configured for server monitoring is **Connected**.

Map IP Addresses to User Names Using Captive Portal

If the firewall receives a request from a zone that has User-ID enabled and the source IP address does not have any user data associated with it yet, it checks its Captive Portal policy for a match to determine whether to perform authentication. This is useful in environments where you have clients that are not logged in to your domain servers, such as Linux clients. This user mapping method is only triggered for web traffic (HTTP or HTTPS) that matches a security rule/policy, but that has not been mapped using a different method.

Captive Portal Authentication Methods

Captive Portal uses the following methods to obtain user data from the client when a request matches a Captive Portal policy:

Authentication Method	Description
NTLM Authentication	The firewall uses an encrypted challenge-response mechanism to obtain the user's credentials from the browser. When configured properly, the browser will provide the credentials to the firewall transparently without prompting the user, but will display a prompt for credentials if necessary. If the browser cannot perform NTLM or if NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration. By default, IE supports NTLM. Firefox and Chrome can be configured to use it. You cannot use NTLM to authenticate non-Windows clients.
Web Form	Requests are redirected to a web form for authentication. You can configure Captive Portal to use a local user database, RADIUS, LDAP, or Kerberos to authenticate users. Although users will always be prompted for credentials, this authentication method works with all browsers and operating systems.
Client Certificate Authentication	Prompts the browser to present a valid client certificate for authenticating the user. To use this method, you must provision client certificates on each user system and install the trusted CA certificate used to issue those certificates on the firewall. This is the only authentication method that enables transparent authentication for Mac OS and Linux clients.

Captive Portal Modes

The Captive Portal mode defines how web requests are captured for authentication:

Mode	Description
Transparent	The firewall intercepts the browser traffic per the Captive Portal rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser will display a certificate error to users attempting to access a secure site. Therefore you should only use this mode when absolutely necessary, such as in Layer 2 or virtual wire deployments.
Redirect	The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the time outs expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they will not need to re-authenticate upon IP address change as long as the session stays open. In addition, if you plan to use NTLM authentication, you must use Redirect mode because the browser will only provide credentials to trusted sites.

Configure Captive Portal

The following procedure shows how to configure Captive Portal using the on-device User-ID agent to redirect requests that match a Captive Portal policy to a Layer 3 interface on the firewall. For instructions on configuring Captive Portal using the Windows-based agent, refer to refer to "Configure User Mapping Using the Windows User-ID Agent" in the PAN-OS Administrator's Guide.

If you plan to use Captive Portal without using the other User-ID functions (user mapping and group mapping), you do not need to configure an agent.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent		
Step 1	Make sure the firewall has a route to the servers it will be monitoring to gather user data (for example, your Domain Controllers and your Exchange servers).	In this release of the product, the firewall must be able to communicate with the servers over the MGT interface, so you must make sure that the network your directory servers are on is accessible from this interface. If this configuration does not work in your environment, you must configure Captive Portal using the Window-based User-ID agent.
Step 2	Make sure DNS is configured to resolve your Domain Controller addresses.	To verify proper resolution, ping the server FQDN. For example: admin@PA-200> ping host dc1.acme.com
Step 3	(Redirect mode only) Create a Layer 3 interface to which to redirect Captive Portal requests.	1. Create a management profile to enable the interface to display Captive Portal response pages: a. Select **Network > Network Profiles > Interface Mgmt** and click **Add**. b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**. 2. Create the Layer 3 interface. See Set Up Interfaces and Zones for instructions. Be sure to attach the management profile you created in Step 1 (on the **Advanced > Other Info** tab of the **Ethernet Interface** dialog). 3. Create a DNS "A" record that maps the IP address you configured on the Layer 3 interface to an intranet host name (that is, a hostname that does not have a dot in the name, such as `ntlmhost`).

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)	
Step 4 (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting requests.You can either generate a self-signed certificate or import a certificate that is signed by an external CA. **Note** When setting up Captive Portal for the first time, imported certificates may not work. If you plan to use an imported certificate, complete the initial configuration without specifying a **Server Certificate**. After you get Captive Portal working, you can go back and switch to the imported certificate.	To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for Captive Portal as follows: 1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **OK** to generate the certificate. 2. To create the certificate to use for Captive Portal, click **Generate**. Enter a **Certificate Name** and enter the DNS name of the intranet host for the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting requests. Click **OK** to generate the certificate. 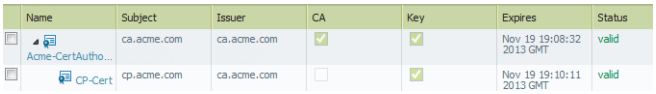 3. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO).
Step 5 Set up an authentication mechanism to use to when the web form is invoked. Note that even if you plan to use NTLM, you must also set up a secondary authentication mechanism that can be used if NTLM authentication fails or if the user agent does not support it. Best Practices: • If using RADIUS to authenticate users from the web form, be sure to enter a RADIUS domain. This will be used as the default domain if users don't supply one upon login. • If using AD to authenticate users from the web form, make sure to enter **sAMAccountName** as the **LogonAttribute**.	1. Configure the firewall to connect to the authentication service you plan to use so that it can access the authentication credentials. • If you plan to authenticate using LDAP, Kerberos, or RADIUS you must create a server profile that instructs the firewall how to connect to the service and access the authentication credentials for your users. Select **Device > Server Profiles** and add a new profile for the specific service you will be accessing. • If you plan to use local database authentication, you must first create the local database. Select **Device > Local User Database** and add the users and groups to be authenticated. 2. Create an authentication profile that references the server profile or local user database you just created. Select **Device > Authentication Profile** and add a new profile for use with Captive Portal.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)	
Step 6 (Optional) Set up client certificate authentication. Note that you do not need to set up both an authentication profile and a client certificate profile to enable Captive Portal. If you configure both, the user will be required to authenticate using both methods. **Note** For details on other certificate profile fields, such as whether to use CRL or OCSP, click the help icon on the web interface page **Device > Certificate Management > Certificate Profile**.	1. Generate certificates for each user that will be authenticating using Captive Portal. 2. Download the CA certificate in Base64 format. 3. Import the root CA certificate from the CA that generated the client certificates onto the firewall: a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**. b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate. c. **Browse** to the **Certificate File** you downloaded from the CA. d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**. e. Select the certificate you just imported on the **Device Certificates** tab to open it. f. Select **Trusted Root CA** and then click **OK**. 4. Create the client certificate profile that you will use when you configure Captive Portal. a. Select **Device > Certificates > Certificate Management > Certificate Profile** and click **Add** and enter a profile **Name**. b. In the **Username Field** drop-down, select the certificate field that contains the user's identity information. c. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in Step 3 and then click **OK**.
Step 7 Enable NTLM authentication. **Note** When using the on-device User-ID agent, the firewall must be able to successfully resolve the DNS name of your Domain Controller in order for the firewall to join the domain. The credentials you supply here will be used to join the firewall to the domain upon successful DNS resolution.	1. Select **Device > User Identification > User Mapping** and click the Edit icon in the **Palo Alto Networks User ID Agent Setup** section of the screen. 2. On the NTLM tab, select the **Enable NTLM authentication processing** check box. 3. Enter the NTLM domain against which the User-ID agent on the firewall should check NTLM credentials. 4. Enter the user name and password for the Active Directory account you created in Step 1 in Map Users to Groups for NTLM authentication.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)

Step 8 Configure the Captive Portal settings.

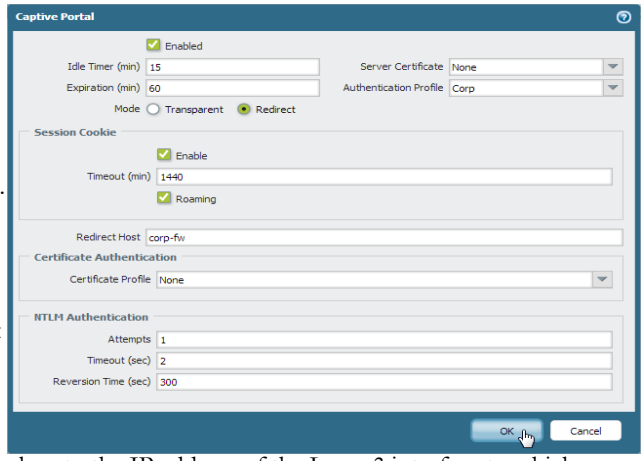

1. Select **Device > User Identification > Captive Portal Settings** and click the Edit icon in the **Captive Portal** section of the screen.
2. Make sure the **Enabled** check box is selected.
3. Set the **Mode**. This example, shows how to set up **Redirect** mode.
4. (Redirect mode only) Select the **Server Certificate** the firewall should use to redirect requests over SSL. This is the certificate you created in Step 4.
5. (Redirect mode only) Specify the **Redriect Host**, which is the intranet hostname that resolves to the IP address of the Layer 3 interface to which you are redirecting requests, as specified in Step 3.
6. (Redirect mode only) **Enable** Session Cookies, using the **Timeout** field to specify the number of minutes the session cookie is valid and select **Roaming** to retain the cookie if the IP address changes while the session is active.
7. Select the authentication method to use if NTLM fails (or if you are not using NTLM):
   - If you are using LDAP, Kerberos, RADIUS, or local database authentication, select the **Authentication Profile** you created in Step 5.
   - If you are using client certificate authentication, select the **Certificate Profile** you created in Step 6.
8. Click **OK** to save your settings.
9. Click **Commit** to save the Captive Portal configuration.

Enable User- and Group-Based Policy

In order to enable security policy based on user and/or group, you must enable User-ID for each zone that contains users you want to identify. You can then define policies that allow or deny traffic based on user name or group membership. Additionally, you can create Captive Portal policies to enable identification for IP addresses that do not yet have any user data associated with them.

Enable User- and Group-Based Policy

Step 1 Enable User-ID on the source zones that contain the users that will be sending requests that require user-based access controls.

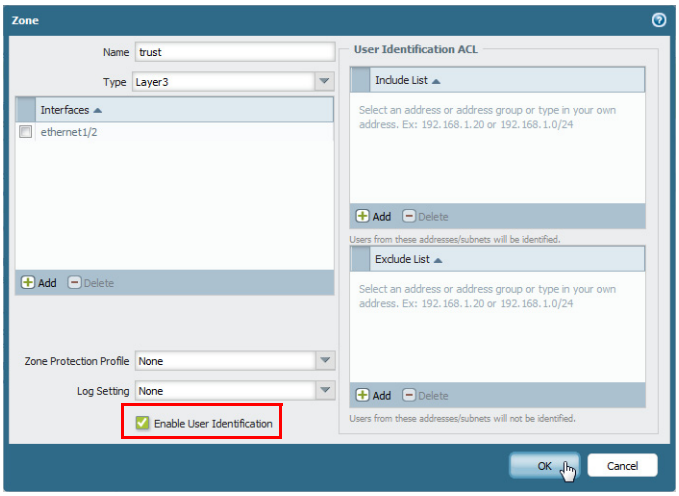
Zone
Name trust
Type Layer3
Interfaces
ethernet1/2
Add Delete
Zone Protection Profile None
Log Setting None
Enable User Identification
User Identification ACL
Include List
Select an address or address group or type in your own address. Ex: 192.168.1.20 or 192.168.1.0/24
Add Delete
Users from these addresses/subnets will be identified.
Exclude List
Select an address or address group or type in your own address. Ex: 192.168.1.20 or 192.168.1.0/24
Add Delete
Users from these addresses/subnets will not be identified.
OK
Cancel

1. Select **Network > Zones**.
2. Click on the **Name** of the zone in which you want to enable User-ID to open the Zone dialog.
3. Select the **Enable User Identification** check box and then click **OK**.

Step 2 Create security policies based on user and/or group.

Best Practice:

Create policy based on group rather than user whenever possible. This prevents you from having to continually update your policies (which requires a commit) whenever your user base changes.

1. After configuring User-ID, you will be able to choose a user name or group name when defining the source or destination of a security rule:
   a. Select **Policies > Security** and click **Add** to create a new policy or click on an existing policy rule name to open the **Security Policy Rule** dialog.
   b. Specify which users and/or groups to match in the policy in one of the following ways:
      – If you want to specify specific users/groups as matching criteria, select the **User** tab and click the **Add** button in the **Source User** section of the dialog to display a list of users and groups discovered by the firewall group mapping function. Select the users and/or groups to add to the policy.
      – If you want the policy to match any user who has or has not successfully authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down list above the **Source User** list.
2. Configure the rest of the policy as appropriate and then click **OK** to save it. For details on other fields in the security policy, see Set Up Basic Security Policies.

Enable User- and Group-Based Policy (Continued)

Step 3 Create your Captive Portal Policies.

1. Select **Policies > Captive Portal**.
2. Click **Add** and enter a **Name** for the policy.
3. Define the matching criteria for the rule by completing the **Source**, **Destination**, and **Service/URL Category** tabs as appropriate to match the traffic you want to authenticate. The matching criteria on these tabs is the same as the criteria you define when creating a security policy. See Set Up Basic Security Policies for details.
4. Define the **Action** to take on traffic that matches the rule. You can choose:
   - **no-captive-portal**—Allow traffic to pass without presenting a Captive Portal page for authentication.
   - **web-form**—Present a Captive Portal page for the user to explicitly enter authentication credentials or use client certificate authentication.
   - **browser-challenge**—Open an NTLM authentication request to the user's web browser. The web browser will respond using the user's current login credentials. If the login credentials are not available, the user will be prompted to supply them.

   The following example shows a Captive Portal policy that instructs the firewall to present a web form to authenticate unknown users who send HTTP requests from the trust zone to the untrust zone.

		Source		Destination			
Name	Tag	Zone	Address	Zone	Address	Service	Action
all	none	trust	any	untrust	any	service-http	web-form

Step 4 Save your policy settings. Click **Commit**.

Verify the User-ID Configuration

After you configure User Identification and enable User-ID on your security policies and Captive Portal policies, you should verify that it is working properly.

Verify the User Identification Configuration	
Step 1 Verify that group mapping is working.	From the CLI, enter the following command: `admin@PA-200>show user group-mapping statistics`

Step 2 Verify that user mapping is working.	If you are using the on-device User-ID agent, you can verify this from the CLI using the following command:
	`admin@PA-200>show user ip-user-mapping-mp all`
	
	`IP Vsys From User Timeout (sec)`
	`------------------------------------------------------------`
	`192.168.201.1 vsys1 UIA acme\louis 210`
	`192.168.201.11 vsys1 UIA acme\eileen 210`
	`192.168.201.50 vsys1 UIA acme\kimberly 210`
	`192.168.201.10 vsys1 UIA acme\administrator 210`
	`192.168.201.100 vsys1 AD acme\administrator 748`
	`Total: 5 users`
	`*: WMI probe succeeded`

Step 3 Test your security policy.	• From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you have defined in your policy and ensure that traffic is being allowed and denied as expected.
	• You can also use the `test security-policy-match` command to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks user duane from playing World of Warcraft, you could test the policy as follows:
	
	`test security-policy-match application worldofwarcraft source-user acme\svogt source any destination any destination-port any protocol 6`
	
	`"deny worldofwarcraft" {`
	`from corporate;`
	`source any;`
	`source-region any;`
	`to internet;`
	`destination any;`
	`destination-region any;`
	`user acme\duane;`
	`category any;`
	`application/service worldofwarcraft;`
	`action deny;`
	`terminal no;`
	`}`

Verify the User Identification Configuration (Continued)

Step 4 Test your Captive Portal configuration.

1. From the same zone, go to a machine that is not a member of your directory—such as a Mac OS system—and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches a Captive Portal policy you have defined. You should see the Captive Portal web form.

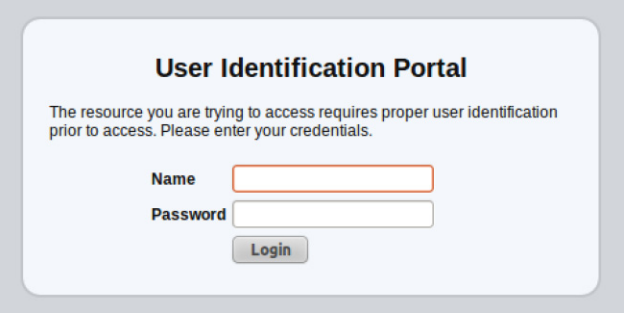
User Identification Portal
The resource you are trying to access requires proper user identification prior to access. Please enter your credentials.
Name
Password
Login

3. Log in using the correct credentials and confirm that you are redirected to the requested page.
4. You can also test your Captive Portal policy using the `test cp-policy-match` command as follows:

```
test cp-policy-match from corporate to internet source
  192.168.201.10 destination 8.8.8.8


Matched rule: 'captive portal' action: web-form
```

Step 5 Verify that user names are displayed in the log files (**Monitor > Logs**).

Receive Time	Category	URL	From Zone	To Zone	Source	Source User
12/18 15:16:17	computer-and-internet-info	*.urbanairship.com/	l3-trust	l3-untrust	10.31.32.18	acme\jdonaldson
12/18 15:16:01	social-networking	graph.facebook.com/	l3-trust	l3-untrust	10.31.32.18	acme\gbalfour
12/18 15:16:01	social-networking	graph.facebook.com/	l3-trust	l3-untrust	10.31.32.18	acme\jdonaldson
12/18 15:16:01	social-networking	orcart.facebook.com/	l3-trust	l3-untrust	10.31.32.18	acme\jparker
12/18 15:04:15	web-based-email	*.mg.mail.yahoo.com/	l3-trust	l3-untrust	10.31.32.18	acme\jreddick
12/18 14:51:13	search-engines	www.google.com/	l3-trust	l3-untrust	10.31.32.18	acme\jparker
12/18 14:49:06	search-engines	*.google.com/	l3-trust	l3-untrust	10.31.32.18	acme\jreddick
12/18 14:48:10	search-engines	www.google.com/	l3-trust	l3-untrust	10.31.32.18	acme\gbalfour
12/18 14:43:53	web-based-email	*.mg.mail.yahoo.com/	l3-trust	l3-untrust	10.31.32.18	acme\jreddick
12/18 14:43:51	internet-portals	android.register.push.mobile.yahoo.com/	l3-trust	l3-untrust	10.31.32.18	acme\jparker
12/18 14:43:51	web-based-email	*.mg.mail.yahoo.com/	l3-trust	l3-untrust	10.31.32.18	acme\jdonaldson
12/18 14:43:51	internet-portals	login.yahoo.com/	l3-trust	l3-untrust	10.31.32.18	acme\jparker
12/18 14:43:50	computer-and-internet-info	*.crittercism.com/	l3-trust	l3-untrust	10.31.32.18	acme\jreddick
12/18 14:36:35	computer-and-internet-info	mdmbeta.paloaltonetworks.com/	l3-trust	l3-untrust	10.31.32.18	acme\gbalfour
12/18 14:36:35	internet-portals	android.connector.push.mobile.yahoo.com/	l3-trust	l3-untrust	10.31.32.18	acme\jdonaldson
12/18 14:36:35	computer-and-	settings.crashlytics.com/	l3-trust	l3-untrust	10.31.32.18	acme\jparker

Verify the User Identification Configuration (Continued)

Step 6 Verify that user names are displayed in reports (**Monitor > Reports**). For example, when drilling down into the denied applications report, you should be able to see a list of the users who attempted to access the applications as in the following example.

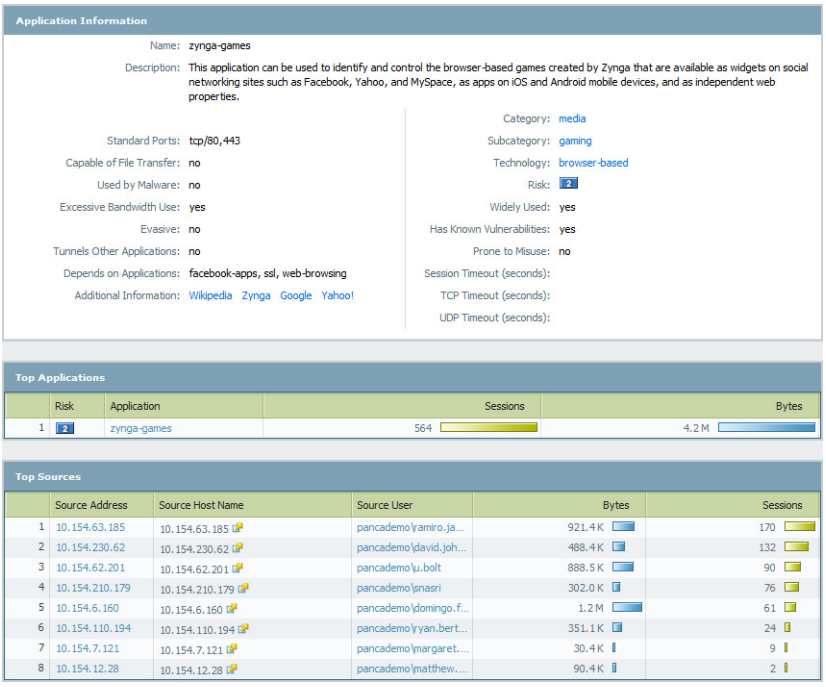

Application Information

Name: zynga-games

Description: This application can be used to identify and control the browser-based games created by Zynga that are available as widgets on social networking sites such as Facebook, Yahoo, and MySpace, as apps on iOS and Android mobile devices, and as independent web properties.

Standard Ports:	tcp/80,443	Category:	media
Capable of File Transfer:	no	Subcategory:	gaming
Used by Malware:	no	Technology:	browser-based
Excessive Bandwidth Use:	yes	Risk:	2
Evasive:	no	Widely Used:	yes
Tunnels Other Applications:	no	Has Known Vulnerabilities:	yes
Depends on Applications:	facebook-apps, ssl, web-browsing	Prone to Misuse:	no
Additional Information:	Wikipedia Zynga Google Yahoo!	Session Timeout (seconds):	
		TCP Timeout (seconds):	
		UDP Timeout (seconds):	

Top Applications

	Risk	Application	Sessions	Bytes
1	2	zynga-games	564	4.2 M

Top Sources

	Source Address	Source Host Name	Source User	Bytes	Sessions
1	10.154.63.185	10.154.63.185	pancademo\ramiro.ja...	921.4 K	170
2	10.154.230.62	10.154.230.62	pancademo\david.joh...	488.4 K	132
3	10.154.62.201	10.154.62.201	pancademo\j.bolt	888.5 K	90
4	10.154.210.179	10.154.210.179	pancademo\snasri	302.0 K	76
5	10.154.6.160	10.154.6.160	pancademo\domingo.f...	1.2 M	61
6	10.154.110.194	10.154.110.194	pancademo\ryan.bert...	351.1 K	24
7	10.154.7.121	10.154.7.121	pancademo\margaret....	30.4 K	9
8	10.154.12.28	10.154.12.28	pancademo\matthew....	90.4 K	2

Set Up High Availability

High availability (HA) is a configuration in which two firewalls are placed in a group to prevent a single point to failure on your network. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity. This section covers the following topics:

- HA Overview
- Prerequisites for Active/Passive HA
- Configuration Guidelines
- Configure an Active/Passive Pair
- Define the Failover Conditions
- Verify Failover

HA Overview

On Palo Alto Networks firewalls, you can set up two devices as an HA pair. HA allows you to minimize downtime by making sure that an alternate device is available in the event that the primary device fails. The devices use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Device specific configuration such as management port IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between devices. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on the active device and the passive device takes over the task of securing traffic, the event is called a failover. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. (Link Monitoring)
- One or more of the destinations specified on the device cannot be reached. (Path Monitoring)
- The device does not respond to heartbeat polls. (Heartbeat Polling)

HA Modes

You can set up the firewalls for HA in two modes:

- **Active/Passive**— One device actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this configuration, both devices share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active device fails, the passive device takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2 and Layer 3 deployments. For information on setting up your devices in an active/passive configuration, see Configure an Active/Passive Pair.

The PA-200 and the VM-Series firewalls support a *lite* version of active/passive HA. HA lite provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization, and therefore, HA Lite does not offer stateful failover.

Active/Active— Both the devices in the pair are active and processing traffic, and work synchronously to handle session setup and session ownership. The active/active deployment is supported in virtual wire and Layer 3 deployments, and is only recommended for networks with asymmetric routing. For information on setting up the devices in an active/active configuration, refer to the Active/Active High Availability Tech Note.

HA Links and Backup Links

The devices in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2)—while others require you to use the in-band ports as HA links.

On devices with dedicated HA ports (HA1 and HA2) such as the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050 firewalls, the dedicated HA ports allow for a direct connection between the management plane and dataplane of the two HA devices. Use these dedicated ports to manage communication and synchronization between the devices. For devices without dedicated HA ports such as the PA-200, PA-500, and PA-2000 Series firewalls, use the management port for the HA1 link, to allow for a direct connection between the management planes on the devices, and an in-band port for the HA2 link.

- **Control Link**: The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. This link is also used to synchronize configuration changes on either the active or passive device with its peer. The HA1 link is a Layer 3 link and requires an IP address.
  Ports used for HA1: TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

- **Data Link**: The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between devices in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device to the passive device. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.
  Ports used for HA2: The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.

  Note: Active/Active deployments also use an HA3 link for packet forwarding.

- **Backup Links**: Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:
  - The IP addresses of the primary and backup HA links must not overlap each other.
  - HA backup links must be on a different subnet than the primary HA links.
  - HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.

Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

Device Priority and Preemption

The devices in an HA pair can be assigned a *device priority* value to indicate a preference for which device should assume the active role and manage traffic. If you need to use a specific device in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each device. The device with the lower numerical value, and therefore *higher priority*, is designated as active and manages all traffic on the network. The other device is in a passive state, and synchronizes configuration and state information with the active device so that it is ready to transition to an active state should a failure occur.

By default, preemption is disabled on the firewalls and must be enabled on both devices. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

Failover Triggers

When a failure occurs on the active device and the passive device takes over the task of securing traffic, the event is called a failover. A failover is triggered when a monitored metric on the active device fails. The metrics that are monitored for detecting a device failure are:

- **Heartbeat Polling and Hello messages**—The firewalls use hello message and heartbeats to verify that the peer device is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the device. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the devices are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. For details on the HA timers that trigger a failover, see HA Timers.
- **Link Monitoring**—The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A device failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link in the link group will cause the device to change the HA state to non-functional to indicate a failure of a monitored object.
- **Path Monitoring**—Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a device failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the device to change the HA state to non-functional to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the device is a suspended state or if preemption occurs.

On the PA-3000 Series, PA-5000 Series, and PA-7050 firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to verify the operational status for all the components within the firewall.

HA Timers

High Availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values across the different hardware models; these values are for current reference only and can change in a subsequent release.

Timers	Description	PA-7050 PA-5000 Series PA-4000 Series PA-3000 Series	PA-2000 Series PA-500 PA-200 VM-Series	Panorama VM M-100
		Current Recommended/Aggressive Values by Platform		
Monitor fail hold up time	The interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices.	0/0	0/0	0/0
Preemption hold time	Time a passive or active-secondary device will wait before taking over as the active or active-primary device.	1/1	1/1	1/1
Heartbeat interval	The frequency at which the HA peers exchange heartbeat messages in the form of an ICMP ping.	1000/1000	2000/1000	2000/1000
Promotion hold time	Time that the passive device (in active/passive mode) or the active-secondary device (in active/active mode) will wait before taking over as the active or active-primary device after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500

Timers	Description	PA-7050 PA-5000 Series PA-4000 Series PA-3000 Series	PA-2000 Series PA-500 PA-200 VM-Series	Panorama VM M-100
		Current Recommended/Aggressive Values by Platform		
Additional master hold up time	This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active device in active/passive mode and to the active-primary device in active/active mode. This timer is recommended to avoid a failover when both devices experience the same link/path monitor failure simultaneously.	500/500	500/500	7000/5000
Hello interval	The time interval in milliseconds between the hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.	8000/8000	8000/8000	8000/8000
Maximum no. of flaps	A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3).	3/3	3/3	Not Applicable

Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—both the devices in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—both the devices should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases. They must also both have the same multiple virtual systems capability (single or multi vsys).
- **The same type of interfaces**—dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the device pair. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For devices without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both devices. However, because the management ports will not be directly cabled between the devices, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 *only* if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
  - Use a crossover cable to connect the HA ports if the devices are directly connected. If the connection is set up using a switch or a router, use a straight through cable.
- **The same set of licenses**—Licenses are unique to each device and cannot be shared between the devices. Therefore, you must license both devices identically. If both devices do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you perform a factory reset on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary device to the newly introduced device with the clean config.

Configuration Guidelines

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both devices and some independently (non-matching) on each device. These HA settings are not synchronized between the devices. For details on what is/is not synchronized, see HA Synchronization.

To proceed with the instructions on configuring the devices in HA, see Configure an Active/Passive Pair.

The following table lists the settings that you must configure identically on both devices:

Identical Configuration Settings on PeerA and PeerB

• HA must be enabled on both devices.
• Both device must have the same Group ID value. The Group ID value is used to create a virtual MAC address for all the configured interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where
00-1B-17: vendor ID; 00: fixed; xx: HA group ID; yy: interface ID.
When a new active device takes over, Gratuitous ARPs are sent from each of the connected interfaces of the new active member to inform the connected Layer 2 switches of the virtual MAC address' new location.
• If using in-band ports, the interfaces for the HA1 and HA2 links must be set to type HA.
• The HA mode must be set to Active Passive.
• If required, preemption must be enabled on both devices. The device priority value, however, must not be identical.
• If required, encryption on the HA1 link (for communication between the HA peers) must be configured on both devices.
• Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup:
  • HA1: Dedicated HA1 port
HA1 Backup: In-band port
Recommendation: Enable Heartbeat Backup
  • HA1: Dedicated HA1 port
HA1 Backup: Management port
Recommendation: Do not enable Heartbeat Backup
  • HA1: In-band port
HA1 Backup: In-band port
Recommendation: Enable Heartbeat Backup
  • HA1: Management port
HA1 Backup: In-band port
Recommendation: Do not enable Heartbeat Backup

The following table lists the settings that must be configured independently on each device:

Independent Configuration Settings	PeerA	PeerB
Control Link	IP address of the HA1 link configured on this device (PeerA).	IP address of the HA1 link configured on this device (PeerB).
	For devices without dedicated HA ports, use the management port IP address for the control link.	
Data Link The data link information is synchronized between the devices after HA is enabled and the control link is established between the devices.	By default, the HA2 link uses Ethernet/Layer 2. If using a Layer3 connection, configure the IP address for the data link on this device (PeerA).	By default, the HA2 link uses Ethernet/Layer 2. If using a Layer3 connection, configure the IP address for the data link on this device (PeerB).
Device Priority (required, if preemption is enabled)	The device you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active device, keep the default value of 100 and increment the value on PeerB.	If PeerB is passive, set the device priority value to a number larger than that on PeerA. For example, set the value to 110.
Link Monitoring—Monitor one or more physical interfaces that handle vital traffic on this device and define the failure condition.	Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover.	Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover.
Path Monitoring—Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness.	Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic. Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover.	Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count.

Configure an Active/Passive Pair

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.

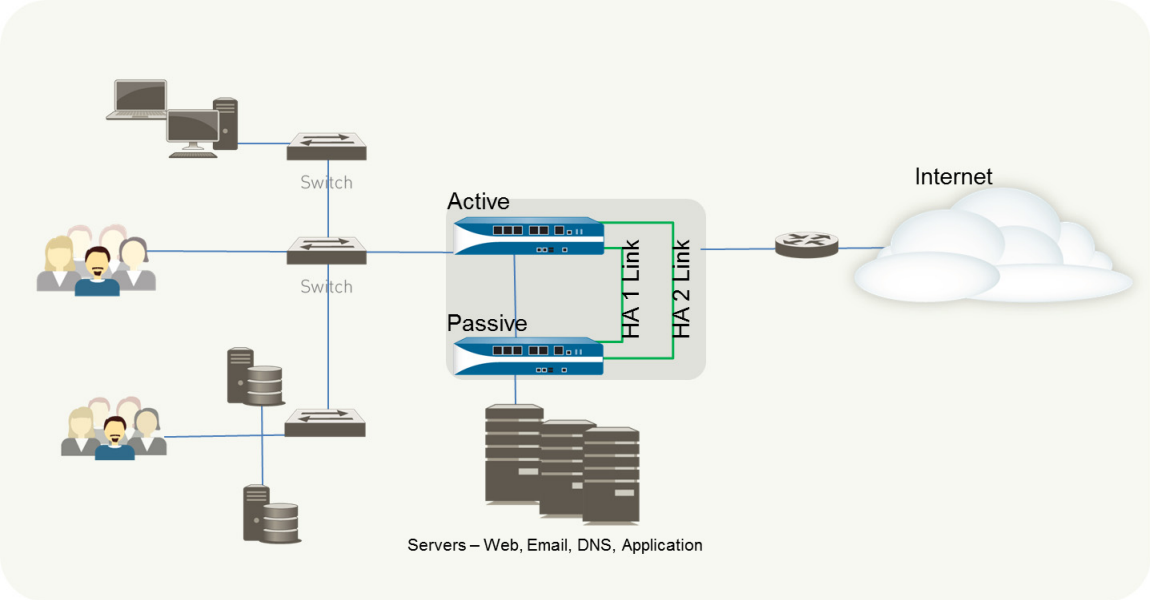

Connect and Configure the Devices	
Step 1 Connect the HA ports to set up a physical connection between the devices.	• For devices with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on the device pair. Use a crossover cable if the devices are directly connected to each other. • For devices without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both devices. Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.
Pick a device in the pair and complete these tasks:	
Step 2 Enable ping on the management port. Enabling ping allows the management port to exchange heartbeat backup information.	1. Select **Device > Setup > Management** and then click the Edit icon in the Management Interface Settings section of the screen. 2. Select **Ping** as a service that is permitted on the interface.
Step 3 If the device does not have dedicated HA ports, set up the data ports to function as HA ports. For devices with dedicated HA ports continue to Step 4.	1. Select **Network > Interfaces.** 2. Confirm that the link is up on the ports that you want to use. 3. Select the interface and set the interface type to HA. 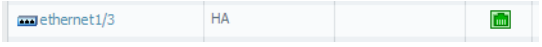 4. On the **Advanced** tab, set the **Link Speed** and **Link Duplex** settings, as appropriate.

Connect and Configure the Devices (Continued)		
Step 4	Set up the control link connection. This example shows an in-band port that is set to interface type HA. For devices that use the management port as the control link, the IP address information is automatically pre-populated.	1. In **Device > High Availability > General**, edit the **Control Link (HA1)** section. 2. Select the interface that you have cabled for use as the HA1 link in the **Port** drop down menu. Set the IP address and netmask. Enter a Gateway IP address only if the HA1 interfaces are on separate subnets. Do not add a gateway if the devices are directly connected.
Step 5	(Optional) Enable encryption for the control link connection. This is typically used to secure the link if the two devices are not directly connected, that is if the ports are connected to a switch or a router.	1. Export the HA key from a device and import it into the peer device. a. Select **Device > Certificate Management > Certificates > Device Certificates.** b. Select **Export HA key.** Save the HA key to a network location that the peer device can access. c. On the peer device, navigate to **Device > Certificate Management > Certificates > Device Certificates,** and select **Import HA key** to browse to the location that you saved the key and import it in to the peer device. 2. Select **Device > High Availability > General**, edit the **Control Link (HA1)** section. 3. Select **Encryption Enabled.**
Step 6	Set up the backup control link connection.	1. In **Device > High Availability > General**, edit the **Control Link (HA1 Backup)** section. 2. Select the HA1 backup interface and set the IP address and netmask. 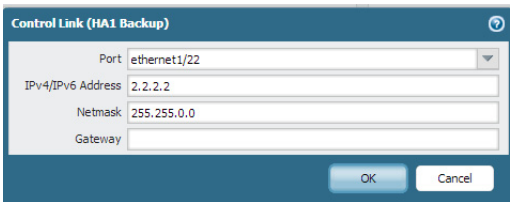

Connect and Configure the Devices (Continued)

Step 7	Set up the data link connection (HA2) and the backup HA2 connection between the devices.	1. In **Device > High Availability > General**, edit the **Data Link (HA2)** section. 2. Select the interface for the data link connection. 3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode. 4. If you use IP or UDP as the transport method, enter the IP address and netmask. 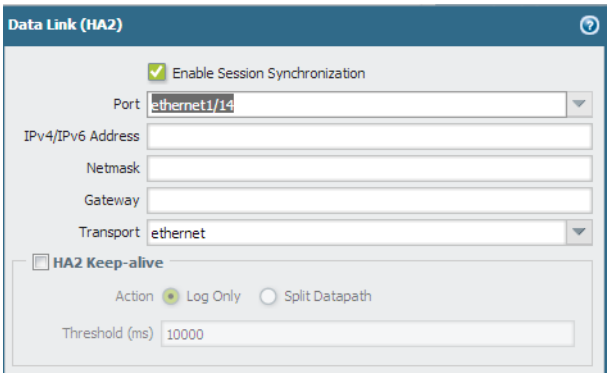 5. Verify that **Enable Session Synchronization** is selected. 6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs. **Note** You can configure the HA2 keep-alive option on both devices, or just one device in the HA pair. If the option is only enabled on one device, only that device will send the keep-alive messages. The other device will be notified if a failure occurs. 7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the IP address and netmask.
Step 8	Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port. You do not need to enable heartbeat backup if you are using the management port for the control link.	1. In **Device > High Availability > General**, edit the **Election Settings** section. 2. Select **Heartbeat Backup**. The heartbeat backup link is used for transmitting redundant heartbeats and hello messages. To allow the heartbeats to be transmitted between the devices, you must verify that the management port across both peers can route to each other.

Connect and Configure the Devices (Continued)	
Step 9 Set the device priority and enable preemption. This setting is only required if you wish to make sure that a specific device is the preferred active device. For information, see Device Priority and Preemption.	1. In **Device > High Availability > General**, edit the **Election Settings** section. 2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the device that you want to assign a higher priority to. **Note** If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active device. 3. Select **Preemptive**. You must enable preemptive on both the active and the passive device.
Step 10 (Optional) Modify the failover timers. By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.	1. In **Device > High Availability > General**, edit the **Election Settings** section. 2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up. **Note** To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen.
Step 11 (Optional, only configured on the passive device) Modify the link status of the HA ports on the passive device. **Note** The passive link state is **shutdown**, by default. After you enable HA, the link state for the HA ports on the active device will be green and those on the passive device will be down and display as red.	Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive device to take over when a failover occurs and it allows you to monitor the link state. To enable the link status on the passive device to stay up and reflect the cabling status on the physical interface: 1. In **Device > High Availability > General**, edit the **Active Passive Settings** section. 2. Set the **Passive Link State** to **Auto**. The auto option decreases the amount of time it takes for the passive device to take over when a failover occurs. **Note** Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered. When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the device.

Connect and Configure the Devices (Continued)	
Step 12 Enable HA.	1. Select **Device > High Availability > General,** edit the **Setup** section. 2. Select **Enable HA.** 3. Set a **Group ID**. This ID uniquely identifies each HA pair on your network, and is essential if you have multiple HA pairs that share the same broadcast domain on your network. 4. Set the mode to **Active Passive**. 5. Select **Enable Config Sync.** This setting enables the synchronization of the configuration settings between the active and the passive device. 6. Enter the IP address assigned to the control link of the peer device in **Peer HA1 IP Address**. 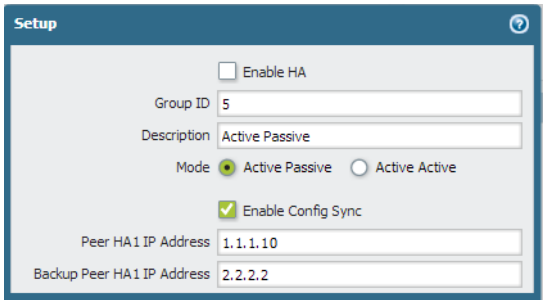 For devices without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer. 7. Enter the **Backup HA1 IP Address**.
Step 13 Save your configuration changes.	Click **Commit**.
Step 14 Complete Step 2 through Step 13 on the other device in the HA pair.	
Step 15 After you finish configuring both devices, verify that the devices are paired in active/passive HA.	1. Access the **Dashboard** on both devices, and view the **High Availability** widget. 2. On the active device, click the **Sync to peer** link. 3. Confirm that the devices are paired and synced, as shown below:

Connect and Configure the Devices (Continued)

On the passive device: The state of the local device should display **passive** and the configuration is **synchronized.**

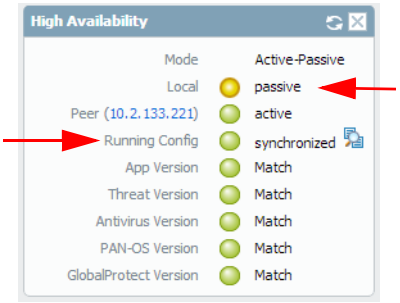

On the active device: The state of the local device should display **active** and the configuration is **synchronized**.

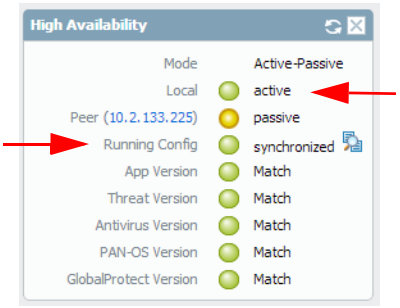

Define the Failover Conditions

Configure the Failover Triggers	
Step 1 To configure link monitoring, define the interfaces that you would like to monitor. A change in the link state of these interface will trigger a failover.	1. Select **Device > High Availability > Link and Path Monitoring.** 2. In the **Link Group** section, click **Add.** 3. Name the **Link Group, Add** the interfaces to monitor, and select the **Failure Condition** for the group. The Link group you define is added to the **Link Group** section. 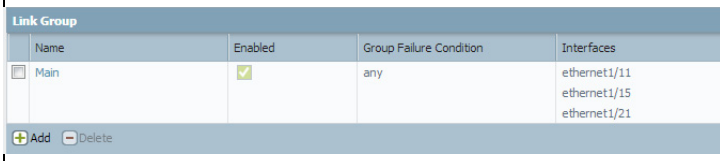
Step 2 (Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the device. By default, the device will trigger a failover when any monitored link fails.	1. Select the **Link Monitoring** section. 2. Set the **Failure Condition** to **All.** The default setting is **Any**.
Step 3 To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity.	1. In the **Path Group** section of the **Device > High Availability > Link and Path Monitoring** tab, pick the option for your setup: **Add Virtual Wire Path**, **Add VLAN Path**, **Add Virtual Router Path**. 2. Select the appropriate item from the drop-down list for the **Name** and **Add** the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the **Failure Condition** for the group. The path group you define is added to the **Path Group** section. 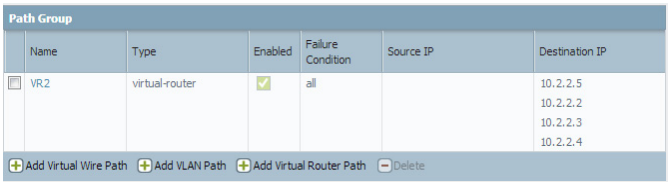
Step 4 (Optional) Modify the failure condition for all Path Groups configured on the device. By default, the device will trigger a failover when any monitored path fails.	Set the **Failure Condition** to **All.** The default setting is **Any**.
Step 5 Save your changes.	Click **Commit**.

Verify Failover

To test that your HA configuration works properly trigger a manual failover and verify that the devices transition states successfully.

Verify Failover	
Step 1 Suspend the active device.	Click the **Suspend local device** link on the **Device > High Availability > Operational Commands** tab. 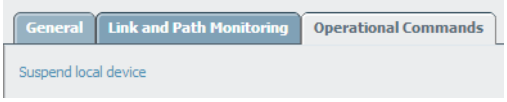
Step 2 Verify that the passive device has taken over as active.	On the **Dashboard**, verify that the state of the passive device changes to **active** in the **High Availability** widget. 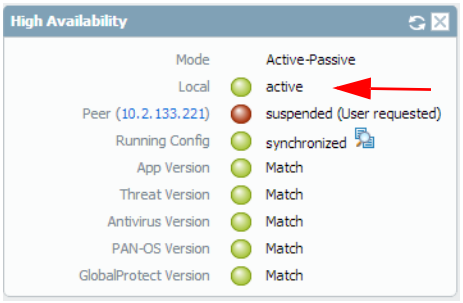
Step 3 Restore the suspended device to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled.	1. On the device you previously suspended, select the **Make local device functional** link on the **Device > High Availability > Operational Commands** tab. 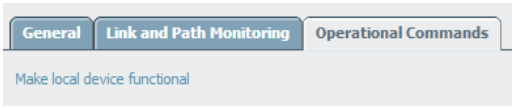 2. In the **High Availability** widget on the **Dashboard,** confirm that the device has taken over as the active device and that the peer is now in a passive state.